

# Partitioning Sparse Graphs using the Second Eigenvector of their Graph Laplacian

David De Wit

June 1991

## Abstract

Partitioning a graph into three pieces, with two of them large and connected, and the third a small “separator” set, is useful for improving the performance of a number of combinatorial algorithms. This is done using the second eigenvector of a matrix defined solely in terms of the incidence matrix, called the graph Laplacian. For sparse graphs, the eigenvector can be efficiently computed using the Lanczos algorithm. This graph partitioning algorithm is extended to provide a complete hierarchical subdivision of the graph. The method has been implemented and numerical results obtained both for simple test problems and for several grid graphs.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	Origins and Acknowledgements . . . . .	4
1.2	Motivation and Applications . . . . .	4
1.3	Broad Outline of the Report . . . . .	7
1.4	Contents of the Report . . . . .	7
<b>2</b>	<b>Theory</b>	<b>8</b>
2.1	M-Matrices . . . . .	8
2.2	Graph Laplacian . . . . .	10
2.3	Graph Partitioning . . . . .	11
2.4	Spectral Partitioning Algorithm . . . . .	15
2.5	Minimum Cover . . . . .	17
2.6	Lanczos Algorithm . . . . .	19
2.7	Recursive Decomposition . . . . .	20
<b>3</b>	<b>Implementation</b>	<b>21</b>
3.1	The Function <code>decomp</code> . . . . .	21
3.2	The Function <code>mk_sp_graph</code> . . . . .	22
3.3	The Function <code>select</code> . . . . .	22
3.4	The Program <code>gr_lap</code> . . . . .	22
3.5	The Program <code>testdc</code> . . . . .	23
<b>4</b>	<b>Results and Problems</b>	<b>24</b>
4.1	Partitioning . . . . .	24
4.2	Full Recursive Decomposition . . . . .	24
4.3	Problems . . . . .	25
<b>5</b>	<b>Further Work</b>	<b>28</b>
<b>6</b>	<b>Listing of Graphic Examples</b>	<b>30</b>
6.1	Idit . . . . .	30
6.2	Moshe . . . . .	32
6.3	Itzchak . . . . .	34
6.4	Shimuel . . . . .	35
6.5	Ishmail . . . . .	38
6.6	Yacov . . . . .	40
<b>7</b>	<b>Code Listing</b>	<b>41</b>
7.1	Listing of <code>decomp</code> . . . . .	41
7.2	Listing of <code>mk_sp_graph</code> . . . . .	49
7.3	Listing of <code>select</code> . . . . .	50
7.4	Listing of <code>gr_lap</code> . . . . .	51
7.5	Listing of <code>testdc</code> . . . . .	52

## List of Figures

1	C Data structure for edge-listing and degree-recording. . . . .	29
2	Idit. . . . .	30
3	The sparse graph Laplacian of Idit. . . . .	31
4	The eigenvalue spectrum of Idit. . . . .	31
5	The second eigenvector of Idit. . . . .	31
6	Moshe. . . . .	32
7	Moshe, after partitioning. . . . .	32
8	The eigenvalue spectrum of Moshe. . . . .	33
9	The second eigenvector of Moshe. . . . .	33
10	The adjacency matrix $L$ of Itzchak. . . . .	34
11	Itzchak. . . . .	35
12	Shimuel. . . . .	35
13	The eigenvalue spectrum of Itzchak. . . . .	36
14	The second eigenvector of Itzchak. . . . .	36
15	The eigenvalue spectrum of Shimuel. . . . .	37
16	The second eigenvector of Shimuel. . . . .	37
17	Ishmail. . . . .	38
18	The eigenvalue spectrum of Ishmail. . . . .	39
19	The second eigenvector of Ishmail. . . . .	39
20	The eigenvalue spectrum of Yacov. . . . .	40
21	The second eigenvector of Yacov. . . . .	40

## List of Tables

1	Times (in CPU seconds) and $ S $ for <b>decomp</b> to partition graphs listed in Table 2. . . . .	25
2	Listing of the experimental graphs. . . . .	30

# 1 Introduction

This report is a detailing of the research carried out by myself during the four months of Semester 1 (March-June) 1991, at the Department of Mathematics, University of Queensland, Australia. It was carried out under the supervision of Dr David E Stewart, and was accredited as a #30 project, subject classification code MN881.

## 1.1 Origins and Acknowledgements

This work was motivated by the publication of a paper [21], titled “Partitioning Sparse Matrices with Eigenvectors of Graphs”, by Alex Pothen, Horst D. Simon and Kang-Pu Liou in the *SIAM Journal on Matrix Analysis and Applications*, September 1990, 11(3):430–452, and emulates some of the implementation contained therein. Their paper is mainly based on some work appearing in an earlier paper [6], titled “A Property of Eigenvectors of Nonnegative Symmetric Matrices and its Application to Graph Theory”, by Miroslav Fiedler in the *Czechoslovak Mathematical Journal*, 1975, 25(100):619–633, which provides a substantial amount of the theory quoted in §2.

I implemented these results in C, interfacing with a library of C data-structures and functions called MESCHACH [23], written by my supervisor, Dr David Stewart. He is also responsible for selected pieces of code that I have used (see Appendix 7), as well as large amounts of time spent educating me in UNIX, C, algorithms, numerical linear algebra and scientific writing. Proof reading was also done by my wife, Nerida Iwasiuk.

## 1.2 Motivation and Applications

The problem dealt with in this report is to partition the set of vertices  $N = \{1, \dots, n\}$  of an undirected, unvaluated graph  $G = (N, E)$  into 3 disjoint sets; a minimally small separator set  $S$ , and 3 “banks”  $A$  and  $B$ , of size of order  $n/2$  each.

Immediate applications of this partitioning apply the “divide-and-conquer” paradigm to a range of graph-theoretic problems such as the travelling sales representative problem. Industrial application problems, in particular those relating to the layout of components in VLSI design, are discussed in [14]. Other application problems are available in the Boeing-Harwell sparse matrix test problem library [3], and include large-scale network-analysis problems such as power-grid distribution problems.

An immediate application in numerical linear algebra is in the efficient solution of large sparse linear systems via factorisation and parallel solution of subproblems. Methods for doing this are outlined below.

Consider the solution of the  $n \times n$  (sparse) symmetric linear system  $C\mathbf{x} = \mathbf{f}$ . If a permutation of the indices of  $C$  can be made, such that in block form:

$$PCP^\top \mathbf{x} = \begin{bmatrix} H_A & \cdot & M_A^\top \\ \cdot & H_B & M_B^\top \\ M_A & M_B & H_S \end{bmatrix} \begin{bmatrix} \mathbf{x}_A \\ \mathbf{x}_B \\ \mathbf{x}_S \end{bmatrix} = \begin{bmatrix} \mathbf{f}_A \\ \mathbf{f}_B \\ \mathbf{f}_S \end{bmatrix},$$

then the system can often be solved much faster. Here  $H_A$ ,  $H_B$  and  $H_S$  are symmetric, and  $M_A$ ,  $M_B$  are usually not (they are usually not even square). The system can be solved using the  $LDL^\top$  (Cholesky-type) factorisation for indefinite, symmetric matrices. Firstly, be aware that simple Cholesky factorisation  $LL^\top$  is only applicable to positive definite matrices, and problems dealt with are **not** guaranteed to be positive definite. (Note that most of the experimental graphs used in this report are 5-point grid graphs, which are always positive semi-definite.) Whilst the graph Laplacian matrices mentioned here are singular, the  $C$  matrices actually involved are not. Writing  $PCP^\top = LDL^\top$  in block form gives:

$$\begin{bmatrix} H_A & \cdot & M_A^\top \\ \cdot & H_B & M_B^\top \\ M_A & M_B & H_S \end{bmatrix} = \begin{bmatrix} L_A & \cdot & \cdot \\ \cdot & L_B & \cdot \\ N_A & N_B & L_S \end{bmatrix} D \begin{bmatrix} L_A^\top & \cdot & N_A^\top \\ \cdot & L_B^\top & N_B^\top \\ \cdot & \cdot & L_S^\top \end{bmatrix},$$

where  $L_A$ ,  $L_B$  are lower triangular,  $L_S$  is symmetric but not generally triangular, and  $N_A$  and  $N_B$  are generally sparse  $|S| \times |A|$  and  $|S| \times |B|$  matrices. In a similar procedure to the usual Cholesky procedure, the solution of the system is done in stages (3, not 2, as there is also a  $D$ ). The first stage in the solution of  $L(DL^\top \mathbf{x}) = \mathbf{f}$  is the solution for  $\mathbf{y} = DL^\top \mathbf{x}$  of  $L\mathbf{y} = \mathbf{f}$  by:

$$\begin{bmatrix} L_A & \cdot & \cdot \\ \cdot & L_B & \cdot \\ N_A & N_B & L_S \end{bmatrix} \begin{bmatrix} \mathbf{y}_A \\ \mathbf{y}_B \\ \mathbf{y}_S \end{bmatrix} = \begin{bmatrix} \mathbf{f}_A \\ \mathbf{f}_B \\ \mathbf{f}_S \end{bmatrix}.$$

To solve this block system, it can be written as:

$$\begin{aligned} L_A \mathbf{y}_A &= \mathbf{f}_A \\ L_B \mathbf{y}_B &= \mathbf{f}_B \\ L_S \mathbf{y}_S &= \mathbf{f}_S - N_A \mathbf{y}_A - N_B \mathbf{y}_B. \end{aligned}$$

The third equation is a (dense) order  $|S|$  system. If  $|S|$  is small, it can be solved at no great expense, using simple Gaussian elimination. This solution depends on the prior solution of the two earlier systems for  $\mathbf{y}_A$  and  $\mathbf{y}_B$ . These are sparse, linear, lower triangular systems of size roughly  $n/2$ , and are cheap to solve using forward substitution.

The second stage is the (trivial) solution of  $D(L^\top \mathbf{x}) = \mathbf{y}$ , using  $L^\top \mathbf{x} = \mathbf{z}$  and  $D\mathbf{z} = \mathbf{y}$ ; which is  $\mathbf{z} = D^{-1}\mathbf{y}$ , as  $D$  is a diagonal matrix. The third stage is similar to the first stage, and provides a solution of  $L^\top \mathbf{x} = \mathbf{z}$ , which is the solution to the whole problem  $C\mathbf{x} = \mathbf{f}$ :

$$\begin{bmatrix} L_A^\top & \cdot & N_A^\top \\ \cdot & L_B^\top & N_B^\top \\ \cdot & \cdot & L_S^\top \end{bmatrix} \begin{bmatrix} \mathbf{x}_A \\ \mathbf{x}_B \\ \mathbf{x}_S \end{bmatrix} = \begin{bmatrix} \mathbf{z}_A \\ \mathbf{z}_B \\ \mathbf{z}_S \end{bmatrix}.$$

Similarly to the first stage, we solve first for  $\mathbf{x}_S$  the order  $|S|$  dense system  $L_S^\top \mathbf{x}_S = \mathbf{z}_S$ , then substitute this into:

$$\begin{aligned} L_A^\top \mathbf{x}_A &= \mathbf{z}_A - N_A^\top \mathbf{x}_S \\ L_B^\top \mathbf{x}_B &= \mathbf{z}_B - N_B^\top \mathbf{x}_S, \end{aligned}$$

to obtain  $\mathbf{x}_A$  and  $\mathbf{x}_B$  cheaply, by backward substitution.

An alternative to the  $LDL^\top$  factorisation is the direct solution of  $(PCP^\top)\mathbf{x} = \mathbf{f}$ . This method involves recursively partitioning the large sparse matrix  $C$ , and is called “Nested Dissection” [8]. We write out the equations in block form:

$$\begin{aligned} H_A \mathbf{x}_A &= \mathbf{f}_A - M_A^\top \mathbf{x}_S \\ H_B \mathbf{x}_B &= \mathbf{f}_B - M_B^\top \mathbf{x}_S \\ H_S \mathbf{x}_S &= \mathbf{f}_S - M_A \mathbf{x}_A - M_B \mathbf{x}_B. \end{aligned}$$

“Solving” the first 2 equations for  $\mathbf{x}_A$  and  $\mathbf{x}_B$ , then substituting the results into the third yields an order  $|S|$  dense system, which gives  $\mathbf{x}_S$  as the solution to:

$$M_A H_A^{-1} \mathbf{f}_A - M_A^\top \mathbf{x}_S + M_B H_B^{-1} \mathbf{f}_B - M_B^\top \mathbf{x}_S + H_S \mathbf{x}_S = \mathbf{f}_S.$$

Explicitly:

$$\mathbf{x}_S = (M_A H_A^{-1} M_A^\top + M_B H_B^{-1} M_B^\top + H_S)^{-1} (\mathbf{f}_S - M_A H_A^{-1} \mathbf{f}_A - M_B H_B^{-1} \mathbf{f}_B).$$

Note that whilst the graph Laplacian is singular, the original  $C$  matrix is generally non-singular, so there is no problem writing this explicitly. After the solution for  $\mathbf{x}_S$ , we have  $\mathbf{x}_A$  and  $\mathbf{x}_B$  as the solution to the following order  $|A|$  and  $|B|$  sparse systems, explicitly written as:

$$\begin{aligned} \mathbf{x}_A &= H_A^{-1} (\mathbf{f}_A - M_A^\top \mathbf{x}_S) \\ \mathbf{x}_B &= H_B^{-1} (\mathbf{f}_B - M_B^\top \mathbf{x}_S). \end{aligned}$$

If  $|A| \approx |B| \approx n/2$ , and  $|S|$  is small, the computational expense is minimised. This situation is amenable to implementation on parallel processing machines, since the 2 linear systems are decoupled. If these are too large to solve simply, the decomposition of  $C$  can be repeated on  $H_A$  and  $H_B$ . The factorisation process can be recursively continued until units of a desired “atomic” size are obtained – perhaps somewhere between 3 and 10. A more thorough analysis could be performed to decide the optimal choice, such that overall computational expense is minimised. This decomposition has been performed for several example graphs in §4.

In this report, the graphs considered are sparse (the average degree of the vertices is low), and large (at current computational capabilities,  $n$  is  $10^3$  to  $10^5$ ). [14, 21] illustrate graphs with  $n$  up to 32400 (for a  $180 \times 180$  5-point grid, average degree 4). An example of a 5-point grid graph is provided in §6.5. (9-point grid graphs are similar, except that each vertex is connected to its nearest 8 neighbours, rather than its nearest 4.) Larger scale applications are available in the Boeing-Harwell sparse matrix test problem library [3], which quotes examples of size up to 44609.

### 1.3 Broad Outline of the Report

The algorithm for partitioning graphs is dependent on analysis of some of the properties of “M-Matrices”, of which the graph Laplacian is an instance. §2.1 describes these entities and some of their properties. The graph Laplacian is introduced, and shown to be an M-Matrix in §2.2.

An amazing result (of Fiedler, 1975 [6]), presented in Algorithm 2.1 on page 15, relates the discrete mathematics of the interconnections of the graph to the continuous mathematics of the eigendecomposition of its “graph Laplacian” (defined in §2.2 in terms of the incidence matrix of the graph). This result takes the components of the second eigenvector of the graph Laplacian matrix, and uses them in a valuation of the **vertices** of the graph. An edge separator set of the graph is then computed, such that the vertices are partitioned by this set into 2 connected banks of very similar size. Using this edge separator set, a combinatorial algorithm creates the 3 vertex sets described in §1.2. This algorithm is fairly efficient, and has a number of immediate applications. Many large tasks can be simplified by using this algorithm in a “divide-and-conquer” approach, which can greatly improve computational efficiency. The Lanczos algorithm (to determine eigenvalues of large, sparse matrices) is implemented as a part of the overall algorithm, and is discussed in §2.6.

This report firstly details experiments performed in synthesising the results of [21], who used the above strategy in computing separator sets of sparse graphs of order  $10^3$  to  $10^5$  vertices. Secondly, it demonstrates the **recursive** decomposition of a graph down to atomic-unit sized subgraphs. This appears to be a new (but probably obvious) implementation.

### 1.4 Contents of the Report

§2 is the largest part of this report, and contains seven subsections. It describes the theory behind the entire graph partitioning process. §3 describes the implementation in C of the algorithms in §2, and references the source code in Appendix 7. §4 mentions the success of the code in dealing with a number of test examples. Results are compared, where possible to literature and touchstone cases (in particular, partitioning 5-point grid graphs). §4 also describes problems encountered in programming and implementation. §5 outlines the directions in which the work could be continued, to make the code more useful. The decomposition of a small example graph (Moshe) is presented in Appendix 6.2.

## 2 Theory

This section has seven subsections:

1. M-Matrices, their origins, applications, definitions, and some interesting theorems.
2. Graph Laplacian, its definition, and properties.
3. Graph Partitioning, using the second eigenvalue of the graph Laplacian [6].
4. Spectral Partitioning Algorithm, developed from the graph partitioning result [21].
5. Minimum Cover – algorithms involved in finding the minimum vertex cover of a bipartite graph.
6. Lanczos Algorithm [9] – required for the implementation of the spectral partitioning algorithm on sparse matrices.
7. Recursive Decomposition – recursive implementation of the partitioning process.

In this report,  $G = (N, E)$  refers to a graph, where  $N = \{1, \dots, n\}$  is the set of vertices.  $A$  always refers to a real, square matrix of size  $n$ . This is a sufficient, but not necessary requirement for some of the definitions and theorems in §2.1. In later sections, specific theorems that also depend on  $A$  being **symmetric** are quoted.

### 2.1 M-Matrices

M-Matrices are a subclass of real matrices that are closely related to non-negative matrices [15], and have a number of interesting properties (for instance, see Theorem 2.5). They arise in several specific fields, e.g.:

1. The discretisation of boundary-value partial differential equations (both symmetric and non-symmetric) generates matrices which are the negative of M-Matrices. Several examples of the discretisation of Poisson's Equation (viz  $\nabla^2 u(x, y) = f(x, y)$ ) on rectangular 5-point grids are explicitly demonstrated in §4 of this report. See also §2.2.
2. Continuous-Time Markov Processes. The differential equations for probabilities are usually (singular) M-Matrices.
3. Economics.

M-Matrices were introduced in 1937 in [17], and their properties have been investigated by a number of researchers. Material presented in this section is abstracted from several sources [2, 6, 15, 19]. M-Matrices have a multiplicity of possible definitions (e.g. [2] lists 50 definitions for nonsingular M-Matrices). This allows great flexibility in proving results which involve them. The following definition of an M-Matrix is quoted from [15].

**Definition 2.1** *A is an **M-Matrix** if there is a non-negative matrix  $B \in \mathbb{R}^{n \times n}$  with maximal eigenvalue  $r$ , and  $c \geq r$  such that  $A = cI_n - B$ .*

The main diagonal entries of an M-Matrix are non-negative and all of its other entries are non-positive. Fiedler [6] calls an M-Matrix a matrix of class  $K_0$ , and a nonsingular M-Matrix a matrix of class  $K$ .

**Definition 2.2**  $F_n = \{A \mid a_{ij} \leq 0, i \neq j\}$ .

That is,  $F_n$  is the set of real matrices whose off-diagonal elements are non-positive. The following results are some necessary and sufficient conditions for an element of  $F_n$  to be an M-Matrix, and are provided mainly as an introduction to some of the properties of M-Matrices. The first is a sufficient condition:

**Theorem 2.1**  *$A \in F_n$  is an M-Matrix iff all its eigenvalues have a non-negative real part.*

The second is a necessary condition:

**Theorem 2.2**  *$A \in F_n$  is an M-Matrix iff every real eigenvalue of  $A$  is non-negative.*

**Definition 2.3** *Given  $M$  as a non-empty subset of  $N$ , define the **principal submatrix** of  $A$  corresponding to  $M$  as  $A(M) = \{a_{ij} \mid i, j \in M\}$ .*

**Theorem 2.3** *A principal submatrix of an M-Matrix is an M-Matrix.*

**Theorem 2.4**  *$A \in F_n$  is an M-Matrix iff all its principal minors are non-negative.*

The next result is important, and is mentioned in [6].

**Theorem 2.5**  *$A \in F_n$  is a nonsingular M-Matrix iff  $A^{-1}$  is non-negative.*

Another definition of M-matrices is:

**Definition 2.4** *A is an **M-Matrix** iff all its off-diagonal entries are  $\leq 0$ , and all its principal minors are non-negative. It is a **nonsingular M-Matrix** if its principal minors are positive.*

Theorems 2.6 and 2.7 are proven in [7] and [24].

**Theorem 2.6** *A nonsingular M-Matrix has  $A^{-1}$  non-negative, and if it is irreducible,  $A^{-1}$  is positive (that is, none of its entries are zero).*

**Theorem 2.7** *If  $A$  is an irreducible, singular M-Matrix, then:*

1. *0 is a simple eigenvalue of  $A$ .*
2. *There is, up to a scale factor, a unique non-zero eigenvector  $\mathbf{u} \in \mathbb{R}^n$  with eigenvalue zero, and all the components of  $\mathbf{u}$  are non-positive or non-negative.*

Theorem 2.8 applies to **symmetric** matrices, and follows from Definition 2.4 and the properties of positive definite (or semidefinite) matrices.

**Theorem 2.8** *A symmetric matrix is an M-Matrix if all its off-diagonal entries are non-positive and all its eigenvalues non-negative. If its eigenvalues are positive, then it is a non-singular M-Matrix.*

## 2.2 Graph Laplacian

This section introduces the graph Laplacian, and defines it in a number of ways.

**Definition 2.5** Given a graph  $G = (N, E)$ , define the **degree** of vertex  $i$  as  $d_i$ , the number of edges in  $E$  with one end being vertex  $i$ .

**Definition 2.6** Given a graph  $G = (N, E)$ , the **graph Laplacian** [1, 6] is the matrix  $L(G)$  (here  $G$  is the incidence matrix of the graph) of the quadratic form:

$$\mathbf{x}^\top L(G) \mathbf{x} = (L(G) \mathbf{x}, \mathbf{x}) = \sum_{(i,j) \in E} (x_i - x_j)^2.$$

$$\text{Thus } L(G) = (l_{ij}), \text{ where } l_{ij} = l_{ji} = \begin{cases} 0 & i \neq j, (i, j) \notin E \\ -1 & i \neq j, (i, j) \in E \\ d_i & i = j. \end{cases}$$

The structure of the graph Laplacian is easier to see in another way. Define  $M \in \mathbb{R}^{n \times n}$  as the **adjacency** matrix of  $G$ , and  $D = \text{diag}(d_1, \dots, d_n) \in \mathbb{R}^{n \times n}$ . The graph Laplacian is then  $L(G) = D - M$ . This definition means that  $L(G)$  is immediately seen as being a matrix with  $-1$  replacing  $1$  in  $M$ , and a diagonal whose  $i$ th element is the number of non-zero off-diagonal elements in row or column  $i$  of  $M$ . This is illustrated by Figures 2 and 3, and is the form used in implementation.

As the row and column sums of the graph Laplacian are zero,  $\mathbf{e} = [1, \dots, 1]^\top \in \mathbb{R}^n$  is an eigenvector corresponding to the eigenvalue 0. As it has a zero eigenvalue, the graph Laplacian is singular. As it satisfies the requirements of Theorem 2.8, it is an M-Matrix. Hence, the graph Laplacian is a singular M-Matrix, and is able to be used in some of the theorems in §2.3.

Experimental graphs considered in this report include several 5-point grid graphs associated with the discretisation of elliptic Boundary-Value PDEs on rectangular regions. If the rectangle is reduced to an  $m \times n$  grid, the graph Laplacian is an  $mn \times mn$  sparse matrix, of very definite structure. Consider **Poisson's Equation** in two dimensions, with boundary conditions for  $u(x, y)$  on a rectangle:

$$\nabla^2 u(x, y) = \frac{\partial^2 u(x, y)}{\partial x^2} + \frac{\partial^2 u(x, y)}{\partial y^2} = f(x, y).$$

(It is **Laplace's Equation** if  $f(x, y) = 0$ .) Discretisation on a regular  $m \times n$  grid, using the notation  $f_{ij} = f(x_i, y_j)$ , and  $u_{ij} = u(x_i, y_j)$  gives:

$$u_{i+1,j} + u_{i-1,j} + u_{i,j-1} + u_{i,j+1} - 4u_{ij} = h^2 f_{ij} \quad \begin{cases} i = 2, \dots, m-1 \\ j = 2, \dots, n-1. \end{cases}$$

This generates an  $mn \times mn$  linear system  $C\mathbf{u} = \mathbf{f}$  on a 5-point grid graph  $G$ . The graph Laplacian  $L(G)$  is an  $m \times m$  block-tridiagonal matrix, where each block is a symmetric  $n \times n$  matrix.

$$L(G) = \begin{bmatrix} P & I_n & . & . & \cdots & . \\ I_n & P & I_n & . & \cdots & . \\ . & I_n & P & I_n & \cdots & . \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ . & . & . & I_n & P & I_n \\ . & . & . & . & I_n & P \end{bmatrix}$$

where:

$$P = \begin{bmatrix} -4 & 1 & . & . & \cdots & . \\ 1 & -4 & 1 & . & \cdots & . \\ . & 1 & -4 & 1 & \cdots & . \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ . & . & . & 1 & -4 & 1 \\ . & . & . & . & 1 & -4 \end{bmatrix}.$$

The structure of  $C$  is closely related to this. A specific example is presented in Appendix 6.5. For 9-point grids or non-elliptic PDEs, the structures are different, but the solution techniques are essentially the same.

## 2.3 Graph Partitioning

This section outlines the work of [6], and theorems listed are taken directly from it.

**Definition 2.7** *Given an undirected, unvaluated graph on  $n$  vertices,  $G = (N, E)$ , a **vertex separator set** is a subset  $M$  of  $N$ , such that removal of the vertices in  $M$  from  $G$ , together with all edges in  $E$  containing them, will disconnect the graph.*

For the purposes of this report, consider only graphs that are connected. We are interested in vertex separator sets that disconnect the graph into 2 subgraphs with approximately equal numbers of vertices. We are particularly concerned with **sparse** graphs and **small** separator sets; but we want algorithms that will automatically perform adequate partitioning of **any** graph.

Ideally, a partitioning algorithm should find a separator set  $S$  of absolute minimal cardinality, but this appears to be a combinatorially explosive problem (probably NP-complete), and is regarded as infeasible. Instead, we attempt to find an  $S$  that is “reasonably” small, at least such that  $|S| \ll n$ , and partitions  $N$  into two sets of about the same size.

It is difficult to quantify good separators [14]. The ideal choice is determined by the particular problem, and it is always possible that a “small” separator set does not **exist**. The extreme example is the completely connected graph, where a separator set is of order of the same size as the original set. For planar graphs, the minimum size of a separator set is always  $\leq \sqrt{8n}$ , and the resulting partition has 2 sides each with no more than  $2/3$  of the total number of vertices [14]. This fact may be of some interest, as many of the large practical problems are either planar, or very nearly so (see [3, 14, 21]). Liu, 1989 [14] describes the notion of “good” separator sets in some detail, and Pothén et al. [21] provide several numerical bounds on the minimum possible sizes of separator sets for a given graph. Many of the bounds may be grossly overestimated, and, as they involve calculation of several of the eigenvalues of the graph Laplacian (an expensive and error-prone procedure), are not particularly useful.

Methods presented in the literature can be found in the references to [3, 11, 13, 14, 21], and are not reiterated here. This section deals with a new process, attributed to [21]. The formal construction of this process depends on theory attributed to [6], which begins with a number of relevant definitions.

**Definition 2.8**  $\text{diameter}(G) = \max_{i,j \in N} \{\text{minimum path length } (i, j)\}$ .

That is, the diameter of a graph  $G = (N, E)$  is the maximum value of the minimum length path between any two vertices.

**Definition 2.9** A *decomposition* of a graph  $G = (N, E)$  is a disjoint vertex cover  $N = N_1 \cup N_2$ , such that  $N_1, N_2 \neq \emptyset$  and  $N_1 \cap N_2 = \emptyset$ .

**Definition 2.10**  $A$  is called *irreducible* if, for no decomposition of  $N = N_1 \cup N_2$ ,  $a_{ij} = 0$ ,  $i \in N_1, j \in N_2$ .

This corresponds to connectedness of graphs. Using Definition 2.10 for symmetric  $A \in \mathbb{R}^{n \times n}$ , define:

**Definition 2.11**  $A$  is of *degree of reducibility*  $d \in [0, n - 1]$  if there exists a decomposition of  $N$  into  $d + 1$  non-empty disjoint subsets  $N = \bigcup_{i=1}^{d+1} N_i$  such that:

1.  $A(N_i)$  are irreducible,  $i = 1, \dots, d + 1$ .
2.  $a_{pq} = 0$ ,  $p \in N_i, q \in N_j, i \neq j$ .

An irreducible symmetric matrix has  $d = 0$ . Thus, a graph has  $d = 0$  if it is connected.

**Definition 2.12** The  $n$  eigenvalues of  $A$  (some are possibly multiple) are ordered in increasing size:  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ .

Similarly, the corresponding eigenvectors are ordered – that is, the  $i$ th eigenvector of  $A$  is the eigenvector corresponding to the  $i$ th smallest eigenvalue. For the purposes of the Graph Partitioning theorem, it will turn out that all eigenvalues are non-negative. The smallest will always be 0, of multiplicity equal to the number of connected components of the graph (or degree of reducibility). For 5-point grid graphs, the largest eigenvalue, by Theorem 2.9 [9, p341], will always be  $\leq 8$ .

**Theorem 2.9 (Gershgorin Circle Theorem)** *Given a matrix  $X^{-1}AX \in \mathbb{C}^{n \times n}$ , such that  $X^{-1}AX = \text{diag}(d_1, \dots, d_n) + F$ , where  $F$  has zero diagonal entries, then:*

$$\lambda(A) \subset \bigcup_{i=1}^n D_i, \quad \text{where } D_i = \left\{ z \in \mathbb{C} \text{ such that } \|z - d_i\| \leq \sum_{j=1}^n |f_{ij}| \right\}.$$

That is, the eigenvalues of a complex matrix lie within the union of a set of  $n$  closed circles in the complex plane. The centre of each circle is at the point corresponding to a diagonal element, and its radius is the sum of the absolute values of the non-diagonal elements of that row. Note that it is usual to use  $X = I$ . As applied to  $A$ , the real graph Laplacian of a grid graph, where:

1. The sums of the absolute values of the non-diagonal entries are exactly equal to the diagonal elements.
2. The diagonal elements of  $A$  are all of value either 2, 3 or 4.
3. All the eigenvalues of  $A$  are real (as  $A$  is real and symmetric).

we obtain  $\lambda(A) \subset [0, 8]$ . More generally, if  $A$  is a graph Laplacian,  $\lambda(A) \subset [0, 2\Delta]$ , where  $\Delta = \max_i d_i$ .

From Theorems 2.7 and 2.8:

**Theorem 2.10** *If  $A$  is symmetric, irreducible, has non-negative off-diagonal entries, and  $A\mathbf{z} = \mathbf{0}$  for some real  $n$ -vector  $\mathbf{z}$ , which is neither zero, positive nor negative; then  $A$  is not positive semidefinite.*

Theorem 2.11 is a corollary to a more general result in [6].

**Theorem 2.11** *If  $A$  is non-negative, irreducible and symmetric, and  $\mathbf{y}_d \in \mathbb{R}^n$  is the  $d$ th eigenvector of  $A$ ,  $d \geq 2$ ; then  $M = \{i \in N | y_i \geq 0\}$  is non-null and the degree of reducibility of  $A(M) \leq d - 2$ .*

This means that, choosing  $d = 2$  and  $\mathbf{v}$  as the second eigenvector of  $A$ ; the degree of reducibility of  $A(M) \leq 0$ , which means that it is 0. Using the second eigenvector of  $A$  as  $\mathbf{v}$  thus yields an irreducible (connected) component.

Fiedler's paper [6] goes beyond the needs of this report in defining the graph Laplacian and associated results for the case of valuated (but not directed) graphs. The following graph-theoretic results are a simplification of those presented in Section 3 of [6].

**Definition 2.13** *A **cut**  $C$  of a graph  $G$  is a set of edges to which a decomposition  $N = N_1 \cup N_2$ , where  $N_1 \cap N_2 = \emptyset$ , exists, such that  $C$  consists exactly of all edges in  $G$  with one vertex in each set of the decomposition. The subgraphs of  $G$  induced by the subsets  $N_1$  and  $N_2$  are called **banks**.*

**Theorem 2.12 (Unique Decomposition of Connected Banks)** *If there is a decomposition  $N = N_1 \cup N_2$  of a graph  $G$ , corresponding to a cut  $C$ , such that both corresponding banks are connected, then the decomposition of  $N$  corresponding to  $C$  is unique.*

**Definition 2.14** The **algebraic connectivity** [5] of  $L(G)$  is defined as the smallest non-zero eigenvalue of  $G$ . Corresponding to this is the **characteristic valuation**, which is an assignment of the elements of the eigenvector corresponding to this eigenvalue of  $L(G)$ .

As mentioned in §2.2, the smallest eigenvalue is always 0, of multiplicity equal to the number of components (connected units) of  $G$ . For connected graphs, the algebraic connectivity and characteristic valuation are equal to the second eigenvalue  $\lambda_2$  and eigenvector  $\mathbf{v}_2$ , respectively.

**Theorem 2.13** For any real  $r$ , define  $M(r) = \{i \in N | y_i \geq -r\}$ . The subgraph  $G(r)$  induced by  $G$  on  $M(r)$  is connected.

Theorem 2.14 is the fundamental result required by the graph partitioning algorithm.

**Theorem 2.14 (Main Graph-Partitioning Theorem)** If there exists a real  $c$  such that  $0 \leq c < \max_i y_i$  and  $c \neq y_i \forall i$ , then the set of edges  $(i, j)$  of  $G$  for which  $y_i < c < y_j$  forms a cut  $C$  of  $G$ . If  $N_1 = \{j \in N | y_j > c\}$  and  $N_2 = \{j \in N | y_j < c\}$ , then  $N = N_1 \cup N_2$  is a decomposition of  $N$  corresponding to  $C$ , and the bank  $G(N_2)$  is connected.

Theorem 2.12 shows that the decomposition and the banks are **uniquely** determined. Define  $N_1 = \{i \in N | y_i > 0\}$  and  $N_2 = \{i \in N | y_i < 0\}$ , then  $N = N_1 \cup N_2$  is the decomposition corresponding to  $C$ . Theorem 2.15, shows that **all** cuts with connected banks in a connected graph are able to be obtained (via the second eigenvector of the graph Laplacian).

**Theorem 2.15** If  $G$  is a connected graph with a cut  $C$  such that both banks of  $C$  are connected then there is a positive valuation of the edges of  $G$  such that the corresponding characteristic valuation  $\mathbf{y}$  is unique (up to a factor), and  $y_i \neq 0 \forall i$ , and  $C$  is formed exactly by alternating edges of the valuation  $\mathbf{y}$ .

In summary, the essential results in [6] are contained in Algorithm 2.1 (that I have composed), which is directly referred to in the beginning of §2.4.

**Algorithm 2.1 (Edge Separator Algorithm)** *Given an graph  $G = (N, E)$ , calculate a edge separator set (cut)  $E_1 \subset E$  such that the 2 resulting banks are connected and have approximately equal numbers of vertices.*

1. Calculate the graph Laplacian  $L(G)$  of  $G$ , and the smallest non-zero eigenvalue of  $G$  (the algebraic connectivity). If  $G$  is connected, the eigenvalue 0 will be of multiplicity 1, in which case, the algebraic connectivity is equal to  $\lambda_2$ , the second eigenvalue of  $L(G)$ .
2. Calculate the corresponding (second) eigenvector (the characteristic valuation).
3. Assign to the vertices of  $G$  the  $n$  elements of the characteristic valuation.
4. Find the set of edges  $E_1 \subset E$ , whose characteristic valuations cross the median value of the components of the second eigenvector.
5.  $E_1$  is the cut required edge separator set. Choice of  $E_1$  from edges whose vertices cross some point between two other components of the characteristic valuation will also yield 2 connected banks, but their sizes will not be nearly equal.

## 2.4 Spectral Partitioning Algorithm

The idea of using the results contained in Algorithm 2.1 in an algorithm to partition graphs into vertex sets appears in Pothen et al., 1990 [21]. They describe their algorithm as a “Spectral Partitioning Algorithm”, and this convention is followed here. They compare the performance of this algorithm with several other algorithms:

1. Kernighan-Lin Algorithm – a modified level structure algorithm that is implemented in SPARSPAK<sup>1</sup>.
2. Fiduccia-Mattheyses Algorithm, implemented by Leiserson and Lewis [13].
3. Separator Algorithm of Liu [14], based on the Multiple Minimum Degree Algorithm.

This report is an emulation of their algorithm, formally described in Algorithm 2.2, largely taken from [21]. This procedure partitions the set of vertices of a (sparse) graph, in the form specified in §1. Firstly, it applies the results in Algorithm 2.1 to the graph concerned, to yield  $E_1$ , an appropriate edge separator set, and  $A'$  and  $B'$ , sets of vertices on either side of this set. Secondly, a combinatorial procedure chooses from the vertices adjacent to this edge separator set, a vertex-separator set  $S$ , and defines the corresponding vertex sets of the banks  $A$  and  $B$ .

Before listing the actual algorithm, the definition of a minimum cover is required.

**Definition 2.15** *Given a graph  $G = (N, E)$ , a (vertex) **cover** is a set of vertices  $S$ , such that every edge in  $E$  has at least one of its endpoints in  $S$ .*

---

<sup>1</sup>SPARSPAK is a package of sparse matrix routines, available through the netlib electronic software library

**Algorithm 2.2 (Spectral Partitioning Algorithm)** *Given the sparse matrix of a graph  $G = (N, E)$ , find a partition of the vertex set  $N = A \cup B \cup S$ , such that  $|A| \approx |B| \approx n/2$ , and  $|S|$  is “small”, in a restricted sense.*

1. *Compute the eigenvector  $\mathbf{x}_2$  and the median value  $x_m$  of its components.*
2. *Partition the vertices of  $G$  into 2 sets,  $A' = \{v \in N \mid x_v \leq x_m\}$  and  $B' = N \setminus A'$ . If  $|A'| - |B'| > 1$ , move enough vertices with components equal to  $x_m$  from  $A'$  to  $B'$  to make this difference at most one.*
3. *Define  $A_1$  as the set of vertices in  $A'$  adjacent to some vertex in  $B'$ , and  $B_1$  as the set of vertices in  $B'$  adjacent to some vertex in  $A'$ . Compute  $H = (A_1, B_1, E_1)$ , the bipartite subgraph induced by the vertex sets  $A_1$  and  $B_1$ .*
4. *The required vertex separator set  $S$  is a minimum cover of  $H$ . It separates  $G$  into subgraphs with vertex sets  $A = A' \setminus S$  and  $B = B' \setminus S$ .*

**Definition 2.16** *A **minimum cover** is a cover of minimum cardinality.*

Associated with the notion of minimum cover is that of a maximum matching, which requires another definition.

**Definition 2.17** *A **matching** is a subset of  $E$ , such that no two endpoints in this subset have the same vertex.*

**Definition 2.18** *A **maximum matching** is a matching of maximal cardinality.*

Maximum matchings and minimum covers are dual concepts, and this is further discussed in §2.5.

Several notes arise in regard to this algorithm:

1. It must be recognised that the ideal aim is to find an  $S$  of absolute minimal cardinality, however this algorithm only finds  $S$  as small as possible in the context of the given edge separator set  $E_1$ . In general, the problem of finding the smallest possible  $S$  appears to be a combinatorially-explosive one that is not achievable by any algorithm efficient enough to be worth considering. In consolation, [21] demonstrate that the Spectral Partitioning algorithm generally finds a smaller  $S$  than its competitors.
2. The problem of finding a minimum cover is non-trivial, and much research into it has been performed. Pothen et al. [21] use a “maximum matching” technique (see §2.5), that is **guaranteed** to give the minimum cover for the given edge separator set. The actual implementation in this report uses a heuristic procedure to calculate an **approximate** minimum cover, described in Algorithm 2.3 in §2.5. This procedure is **not** guaranteed to give a minimum cover.
3. Algorithm 2.2 requires the use of a function to determine the median of a list of numbers. An algorithm to do this efficiently is a special case of an

algorithm to select the  $k$ th smallest component of a list of numbers. Page 129 of [22] describes an algorithm to do this in  $O(n \log n)$  time, by recursively partitioning the list.

4. The problem of calculating the second eigenvector of the graph Laplacian is also non-trivial, and generally computationally expensive. It depends on the prior calculation of the second eigenvalue. The most efficient algorithm to find extremal (largest and smallest) eigenvalues of sparse matrices is the Lanczos algorithm. Unfortunately, the procedure is subject to severe numerical problems that make implementation complex, but in practice it is the only real choice. Implementation of the Lanczos algorithm is discussed in more detail in §2.6.

## 2.5 Minimum Cover

This section deals with the problem of finding a minimum cover (see Definition 2.16) of a **bipartite** graph. This problem occurs as a necessary component of the Spectral Partitioning algorithm (Algorithm 2.2) – it is desired to find a minimum cover of  $H = (A_1, B_1, E_1)$ . One method of doing this is mentioned in [21], but this has not been implemented due to time constraints, and instead a heuristic procedure is followed. §2.5.1 discusses references in which are found minimum cover algorithms (for both bipartite and general graphs), and §2.5.2 describes the heuristic procedure (for bipartite graphs only) actually implemented by me.

### 2.5.1 True Solution

The minimum cover problem has been solved in a number of ways. The earliest algorithms for the minimum cover of a bipartite graph are based on the “Dulmage-Mendelsohn decomposition” [4, 10], but these references are not particularly readable. A more general result, for the minimum cover of any graph is found in [16], but for the purposes of this partitioning, it is better to only consider algorithms for bipartite graphs, to maximise efficiency. For the rest of this section, consider the term graph to mean the bipartite graph  $H = (A_1, B_1, E_1)$ . Also define  $a = |A_1|$ ,  $b = |B_1|$ ,  $e = |E_1|$  and  $n = a + b$ .

As mentioned in §2.4, the minimum cover of a bipartite graph is the dual of the maximum matching, although the details of this relationship are not explicitly supplied. Pothen et al. [21] cite a further paper [20] that details an algorithm for a maximum matching. A simplified description of this algorithm is found in [18, pp221–227], and is reported to solve the matching problem in  $O(\min(a, b).e)$  time.

An alternative approach is presented in pages 495–499 of [22], and deals with the matching problem in terms of a problem in “flow maximisation”. The algorithm described is quoted as requiring  $O(n(e + n) \log n)$  time (or  $O(n^3)$  time, for dense graphs). It is expected that  $H$  is in general sparse, but this algorithm does not appear to be as efficient as that used in [21].

Implementation of a true minimum cover algorithm is left as a future exercise, see §5.

### 2.5.2 Heuristic Solution

The heuristic procedure which I have used is formally described in Algorithm 2.3.

Input data is an edge separator set  $E_1$ , generated by the vertices whose eigenvaluation crosses the median eigenvalue. Construct  $d$ , a listing of the degrees of the vertices in  $H = A_1 \cup B_1$ , and then, whilst any element in  $d$  is positive, perform the following procedure: Find the current vertex of maximum degree in  $d$ , add it to  $S$ , and remove it and the edges adjacent to it (in  $E_1$ ), from  $d$ . Repeat until no edges of degree  $> 0$  remain in  $d$ . Setting  $d_j$  to 0 ensures that the vertex will not be considered again. Implementation is improved by the very cheap process of comparing the size of the resultant  $S$  with the sizes of  $A_1$  and  $B_1$ , and if  $S$  is larger than the smaller of these,  $S$  is replaced by the smaller one. Since no elements are removed from  $S$ , it can be implemented as a simple list or array.

**Algorithm 2.3 (Approximate Minimum Cover Algorithm)** *Given  $E_1$ , the edge separator set of edges of a bipartite graph on vertex sets  $A_1$  and  $B_1$ , find a vertex (separator) set  $S$  such that each edge in  $E_1$  is incident to a vertex in  $S$ , with  $|S|$  “small”.*

```

 $S \leftarrow \emptyset;$ 
for  $i \in A_1 \cup B_1$ 
     $d_i \leftarrow$  degree of  $i$  in  $E_1$ ;
while some  $d_i > 0$ 
    choose  $j$  such that  $d_j = \max_i d_i$ ;
     $d_j \leftarrow 0$ ;
     $S \leftarrow S \cup \{j\}$ ;
    for each  $i$  adjacent to  $j$  in  $E_1$ 
         $d_i \leftarrow d_i - 1$ ;
if  $|S| > |A_1|$ 
     $S \leftarrow A_1$ ;
if  $|S| > |B_1|$ 
     $S \leftarrow B_1$ ;

```

This strategy will work to varying degrees with different examples. In practice, it would be possible to create a (possibly pathological) example for which this algorithm would create an unreasonably large  $S$ . For sparse graphs of large diameter, it is expected that  $|A_1| \ll |A'|$  and  $|B_1| \ll |B'|$ , so the process for finding  $S$  should yield the desired set  $|S| \ll n$ . It is **guaranteed** only that  $|S| \leq \min(|A_1|, |B_1|)$ .

**Algorithm 2.4 (The Lanczos Algorithm)** *Given a symmetric  $A \in \mathbb{R}^{n \times n}$  and  $\mathbf{w} \in \mathbb{R}^n$  having unit 2-norm, compute a symmetric, tridiagonal matrix  $T_j \in \mathbb{R}^{j \times j}$  with the property that  $\lambda(T_j) \subset \lambda(A)$ . The diagonal and subdiagonal elements of  $T_j$  are stored in  $\alpha(1:j)$  and  $\beta(1:j-1)$  respectively.*

```

v = 0   $\beta_0 = 1$ ;   $j = 0$ ;
while  $\beta_j \neq 0$ 
    if  $j \neq 0$ 
        t = w;  w =  $\frac{1}{\beta_j} \mathbf{v}$ ;  v =  $-\beta_j \mathbf{t}$ ;
    end
    v = v +  $A\mathbf{w}$ ;   $j = j + 1$ ;
     $\alpha_j = \mathbf{w}^\top \mathbf{v}$ ;  v = v -  $\alpha_j \mathbf{w}$ ;   $\beta_j = \|\mathbf{v}\|_2$ ;
end

```

## 2.6 Lanczos Algorithm

The Lanczos algorithm (originally attributable to [12]) is an efficient method of finding the extremal eigenvalues of a sparse matrix. As the Spectral Partitioning algorithm is to be applied to sparse matrices, the Lanczos algorithm is required in its implementation. It is listed formally in Algorithm 2.4, copied almost verbatim from Chapter 9 of [9].

Practical implementation of the Lanczos algorithm almost always requires some reorthogonalisation. For the purposes of the Spectral Partitioning algorithm, we need only reorthogonalise against  $\mathbf{e}$  (the  $n$ -vector of ones), as the subspace required must be orthogonal to this. Rounding error can be avoided by the following strategy: When computing  $\mathbf{y} = A\mathbf{x}$ , instead of simply returning  $\mathbf{y}$ , we will return  $\mathbf{y}$  orthogonalised against  $\mathbf{e}$ :

$$\mathbf{y} \leftarrow \mathbf{y} - \frac{\mathbf{e}^\top \mathbf{y}}{n} \mathbf{e}.$$

That is, return  $(I - \mathbf{e}^\top \mathbf{e}/n)A\mathbf{x}$  instead of  $A\mathbf{x}$ , as well as starting with  $\mathbf{x}_0 \perp \mathbf{e}$ . The latter requirement is satisfied by the choice of  $(\mathbf{x}_0)_i = i - (n+1)/2$ ,  $i = 1, \dots, n$ , a choice recommended in [21].

## 2.7 Recursive Decomposition

It is possible, and may be necessary for many applications, to be able to repeat the decomposition process on the subcomponents  $A$  and  $B$  of  $N$ . Once the Spectral Partitioning algorithm is implemented, it can be recursively called, until the subgraphs  $A$  and  $B$  are of a certain manageable “atomic” size (about 3 – 10 vertices), yielding a complete permutation of the vertices in  $N$ , via the procedure described in Algorithm 2.5.

**Algorithm 2.5 (Recursive Decomposition Algorithm)** *Given a graph  $G = (N, E)$ , on the set of vertices  $N = \{1, \dots, n\}$ , decompose  $N$  into a permutation of itself, such that the smallest size units in the permutation are no larger than  $p$ , the “atomic” size.*

**function**  $partition(N)$

*Apply the Spectral Partitioning Algorithm to partition  $N$  into  $A \cup B \cup S$ ;*

**if**  $|A| > p$

$A \leftarrow partition(A)$ ;

**if**  $|B| > p$

$B \leftarrow partition(B)$ ;

**return**  $(A \mid B \mid S)$

### 3 Implementation

The algorithms presented in §2 are implemented as a small collection of functions and programs in ANSI-standard C that are interfaced with a C software library (MESCHACH [23]). They were written within a UNIX (BSD4.3) environment. Appendix 7 provides a listing of the source code, although the functions and data-structures in MESCHACH that are referenced are not explicitly listed, and the reader is referred to [23].

Sections 3.1 to 3.5 describe the operation and application of each of the programs/functions listed in Appendix 7. Currently, there is no documentation on these codes apart from this section and comments in the relevant source code. §4 discusses the results of applying these functions to the sample graphs described in Appendix 6.

#### 3.1 The Function `decomp`

The function `decomp` is the primary partitioning (and recursive decomposition) routine written. Input/output parameters are:

1. `sp_mat *L`: Pointer to the sparse matrix of the graph Laplacian of the graph  $G$  that we wish to partition, of size  $n$ . Unchanged on exit.
2. `PERM *P`: On entry, a pointer to a permutation (list) of length  $n$  of the actual numbers of the vertices in the set being partitioned.  $P$  is returned as the permutation of itself corresponding to the partition.
3. `PERM *A`, `PERM *B`, `PERM *S`: The components of the partition. Currently not actually relevant, but will be used in future developments, where the structure of a recursive decomposition will also be returned. These are pointers to permutations of size  $n$  on entry, containing dummy data. Returned as correctly-sized and filled permutations.
4. `int rec_lvl`: `decomp` is designed perform one of two types of tasks:
  - (a) Partition a sparse graph once.
  - (b) Recursively repeat this until the units involved are of size less than a nominated (currently hard-wired into `decomp` as 3) “atomic” size.  $P$  is returned as the permutation which will be most useful for subsequent factorisation. This is not directly useful, at present, as there is no record of the **structure** of  $P$  – again, this is left as future work, and is discussed in §5.

`rec_lvl` is used to tell `decomp` which of these 2 tasks to perform. If `rec_lvl` is set to -1 by a driver program, `decomp` will only partition the graph once, any other choice allowing it to recursively partition the set until satisfied. The parameter is used as a record of the depth of recursion by `decomp`, so a driver program will typically only ever set `rec_lvl` to 0 or -1.

Currently, `decomp` is called by the driver program `testdc` (see §3.5), and the parameter `rec_lvl` is set by `testdc` according to one of its input (`argv`) arguments.

### 3.2 The Function `mk_sp_graph`

A function to generate random sparse graphs. It is not usually used for testing `decomp`, as the graphs generated are possibly not connected, and are not of small diameter, which means that we cannot expect small separator sets from them. Future versions of this function should generate essentially planar graphs of large diameter, which we could expect to partition into sets with small  $S$ . Input parameters are:

1. `unsigned int n`: The order of the graph desired.
2. `unsigned int p`: The average degree of the vertices.

The function returns a pointer to a `sp_mat`, and is called by `testdc`, a driver program for `decomp`.

### 3.3 The Function `select`

This function selects the  $k$ th smallest element of a series of (real) numbers, stored in a `VEC *`, (a pointer to a `VEC`). It is an implementation of an efficient algorithm in [22, page 128]. In particular, it can be used to determine the median element of a series of numbers.

Input parameters are:

1. `VEC *y`: The vector involved.
2. `int k`: `select` returns the **value** (not the position) of the  $k$ th smallest element of  $y$ .

`select` returns a `double`. `decomp` calls it to find the second smallest element in the vector of “eigenvalues” returned by `trieig`. The function is also called by `decomp` to determine the second smallest eigenvalue in a list (see Algorithm 2.2).

### 3.4 The Program `gr_lap`

This program generates sparse matrices of graphs associated with solving Laplace’s Equation on a rectangular 5-point grid. The graphs, not the graph Laplacians are generated. `gr_lap` automatically creates a file of the appropriate name, which can be changed for later use. For example (in a UNIX environment), typing the command:

```
gr_lap 13 23
```

creates a file called “`Lap.13.23`” in the current directory, which contains (in standard format for MESCHACH to read), the sparse graph related to solving Laplace’s Equation on a  $13 \times 23$  grid.

### 3.5 The Program `testdc`

This program is the main driver written for `decomp`, and is called with a series of (`argv`) input parameters. For example (in a UNIX environment), typing the command:

```
testdc 157 7 13 -1
```

calls `testdc`, and tells it that we wish to use a graph on 157 vertices, with an average vertex degree of 7, and to only print out intermediate results of data-structures of size 13 or less. `testdc` will then prompt the user to select either one of a range of graphs stored in the current directory, or generate a random sparse graph, for partitioning. If the user chooses to input a graph that does not have the correct dimension, `testdc` will abort. This stupid-seeming system is also intended to allow the user to instruct `testdc` to generate random sparse graphs (using `mk_sp_graph`), with parameters `n` = 157, and `p` = 7. Thus, the user is expected to have some knowledge of the database of sparse graphs before using `testdc`. The system will be improved in subsequent versions of `testdc` and `decomp`.

The final parameter is the value of `rec_lvl` (see §3.1) that we wish to initially pass to `decomp`. If set to -1, `decomp` is instructed to only partition the graph once. Any other value, or omission of it instructs `decomp` to recursively decompose the graph into atomic subunits. The atom-size is currently hard-wired into `decomp` (as 3), but could become an input parameter in future versions. `testdc` currently calls `select` and `decomp`, as well as numerous subroutines from `MESCHACH`.

## 4 Results and Problems

This section details some results obtained by

1. Partitioning of all the graphs listed in Table 2 (page 30 in Appendix 6).
2. Full recursive decomposition of the 5 smallest ones.

It refers to the example graphs listed in Appendix 6. These graphs were used in the development of the programs, and have been used to illustrate of the progress of the algorithms. There is only one picture of a grid graph – for Ishmail, on 55 vertices. The other grid graphs are too large to depict. The successful partitioning of these graphs is mentioned in §4.1, recursive decomposition in §4.2, and problems encountered are discussed in §4.3.

Direct comparison with other published results (in particular [21]) is not possible, as the computers involved are of very different speeds. The test problems that have been used are largely from the Boeing-Harwell sparse matrix test problem library [3]. Access paths to this database were discovered too late for it to be used. Other implementational problems, such as the Boeing-Harwell database being stored as a column-oriented data-structure, made the application of the partitioning algorithm to this suite currently infeasible. (MESCHACH deals primarily in terms of row-oriented data-structures.)

### 4.1 Partitioning

**decomp**, driven by **testdc**, has been used to correctly partition all of the graphs listed in Table 2. The times (in CPU seconds) taken to do this on the University of Queensland's Mathematics Department Pyramid 9810 computer (operating under UNIX BSD4.3) are listed in Table 1. The Pyramid 9810 is benchmarked by a LINPACK<sup>2</sup> routine at approximately  $0.5 \text{ Mflop s}^{-1}$ . The computer used by Pothen et al. [21] is a CRAY Y-MP,<sup>3</sup> and is expected to be benchmarked about 2 orders of magnitude faster than the Pyramid 9810. Results are roughly comparable with those in [21], by incorporating a scaling factor of about 100.

All but one of the separator sets listed in Table 1 are of the minimum possible size. The separator set of size 101 for Schlomo could have been improved (to size 61), by changing tolerances used in **decomp**, but this would have required prohibitive amounts of memory.

### 4.2 Full Recursive Decomposition

The full recursive decomposition was successfully completed for the graphs Idit, Moshe, Itzhak, Shimuel, Ishmail, and Yacov, finding permutations of the vertices that were traced to be exactly what would be expected for the first four cases. **decomp** was allowed to run on Ishmail and Yacov until it had completed a full recursive decomposition – this ran to 6 levels of recursion for Yacov. Results appeared correct, although were not manually checked explicitly! The full recursion on the other (very large) problems was not attempted.

---

<sup>2</sup>LINPACK is available through the netlib electronic software library

<sup>3</sup>Trademark of CRAY Research

Name	$n$	$ S $	$S$	CPU time (s)
Idit	7	1	3	0.0
Moshe	20	2	9, 10	0.1
Itzhak	31	1	30	0.2
Shimuel	32	2	7, 11	0.1
Ishmail	55	5	5, 16, 27, 38, 49	0.1
Yacov	105	5	9, 31, 52, 73, 94	0.3
Yair	121	11	59–69	0.2
Arieh	505	5	50, 151, 252, 353, 454	8.1
Aaron	2121	21	50, 151, 252, ..., 2070	22.0
Schlomo	6161	101	3079–3179	44.4
Shimshon	6400	80	3120–3199	57.8

Table 1: Times (in CPU seconds) and  $|S|$  for **decomp** to partition graphs listed in Table 2.

Caveat: the above comments refer to a **previous** version of **decomp**, not the one supplied in Appendix 7, which has some minor implementational problems. Currently, **decomp** crashes at some point during recursive calls.

### 4.3 Problems

This section describes the most significant problem encountered in implementing the Spectral Partitioning algorithm (Algorithm 2.2 on page 16). This problem is in the choice of  $m$ , the order of the tridiagonal matrix  $T_m$  returned by the Lanczos algorithm.

The first impression is that any choice of  $m$  will do, the larger the better, in estimating  $\lambda_2$  accurately. The constraint is that the computational expense is dependent on  $m^2$  (amongst other things), and  $m$  cannot be increased without limit. At the very least,  $m \leq n$  would appear to be an obvious, if exaggerated bound. There is then, an optimal choice of  $m$ , such that:

1.  $\lambda_2$  is found accurately.
2. Minimal computational expense is involved.

In fact, empirical observation demonstrates that there is typically a range of suitable  $m$ , this range varying in position and bandwidth with the problem encountered. It is **not** possible to accurately predict this range in advance, although typically it is located in the vicinity of  $\sqrt{n}$ . ( $n$  is the number of vertices in the graph.)

To make matters more difficult, if  $m$  is chosen from outside this range, the computed value of  $\lambda_2$  will be **wrong**. It appears that if  $m$  is too small,  $\lambda_2$  will be too large, sometimes not even corresponding to larger eigenvalues of the graph Laplacian. If  $m$  is too large, the calculated  $\lambda_2$  cannot possibly be an eigenvalue (as it is supposed to be the second smallest one!), and with increasing  $m$  the returned values of  $\lambda_2$  typically reduce, suddenly hitting 0, and remaining there. This phenomenon is called the “Ghost Eigenvalue” problem [9], and has been studied but apparently not yet conquered.

Thus, a strategy is needed to choose an optimal  $m$ . Empirical observations suggest that if  $m$  is increased from below the optimal position, when the resulting computed values of  $\lambda_2$  converge, they converge to the **correct** value. This is not necessarily true when decreasing  $m$  from above the optimum, but it may be so.

Several approaches are possible in developing a strategy for choosing  $m$ :

1. Begin with a value of  $m$  **much** larger than is expected to be adequate for the given problem. For example, if  $m = 20$  is guessed to be optimum, try  $m = 40$ . Using this value, run the Lanczos algorithm and investigate the resulting  $\lambda_2$ . As the Lanczos algorithm has returned us  $T_m$ , by decrementing  $m$  and considering  $T(1 : m - 1, 1 : m - 1)$ , it is possible to calculate  $\lambda_2$  using  $m - 1$ . This procedure is repeated until the computed  $\lambda_2$  converge. It is not expensive, as no new calculations of  $T_m$  have to be made. The problems with this method are that:
  - (a) It is not certain that convergence from above generates the correct  $E_1$ .
  - (b) Calculation of  $T_m$  for an overly large  $m$  is expensive.
  - (c) The choice of the initial  $m$  is problem-dependent. If this choice cannot be automated, there is no hope of an automatic sparse matrix factorisation technique being derived. The only obvious way to automate this choice would be to use  $m = n$ , and this would be so expensive as to make the whole technique infeasible.

This strategy has been considered, but was not implemented due to these concerns. Its prime advantage is in its ease of programming.

2. The second method is to begin with a small  $m$ , say  $m = \min(n, 3)$ , and always maintain  $m$  at least this size. Run the Lanczos algorithm using this value of  $m$ , and return. First calculate the resulting  $\lambda_2$ , and then the value that  $\lambda_2$  would have been if  $m$  were one less. If these two values are the same, accept  $m$  and  $\lambda_2$ , else increment  $m$  by 2 and re-run the Lanczos algorithm. This technique is suited to calling the Lanczos algorithm as an external function, but suffers from being wasteful in its computational requirements, as the  $T_m$  calculations have to be repeated each time  $m$  is incremented. In practice, this technique has been experimented with, and works, but has since been superseded by the next technique, which is much cheaper to implement, although more difficult to program.
3. Follow a similar strategy to the above, but maintain all the information on the Lanczos algorithm as it progresses. In order to do this, a modified Lanczos algorithm is incorporated into the part of the code that examines the convergence of the computed second eigenvalue, so that the need for external function calls is removed. This technique is incorporated into the current version of **decomp**, that appears in Appendix 7.

If the problem being considered is on  $n$  vertices, with average degree  $k$ , and  $m$  is the actual size of the computed  $T$  matrix, then the computational cost of the second and third methods can be compared by the following analysis. [9] describes the cost of one Lanczos iteration (without reorthogonalisation) as  $(2k + 8)n$  flops.

For the second method, returning  $T_j$  requires  $(2k + 8)nj$  flops. If this cost is repeated for  $j = 3, 5, \dots, m$ , then the total cost is:

$$\sum_{j=3}^m (2k + 8)nj = \sum_{i=1}^{(m-1)/2} (2k + 8)n(2i + 1) \approx (k + 4)nm^2/2 \text{ flops.}$$

For example, for the  $21 \times 101$  5-point grid graph Aaron, where  $k \approx 4$ ,  $n = 2121$  and observation suggests that  $m = 35$  is a good choice, this represents  $\approx 10^7$  flops.

Compare this with the cost for the third method, which is simply the cost of the last Lanczos algorithm call of the second (plus some small overhead). This is approximately  $(2k + 8)mn \approx 2 \times 10^6$  flops, an improvement of order  $m/4$ . This means, for example, using the Pyramid 9810 computer, running at  $\approx 0.5 \times 10^6$  flops  $\text{s}^{-1}$ , calculating the second eigenvector takes 2 rather than 20 CPU seconds.

In summary, none of the methods described are **guaranteed** to work, as they are all based on the observation that when the computed second eigenvalues converge, they converge to the correct value of  $\lambda_2$ . Thus their philosophy is heuristic, and needs refinement, but it **appears** to work for the examples tested. This is an outstanding problem for further work.

## 5 Further Work

This section is a listing of directions for further work, and includes some commentary of the work of [21]:

The incorporation of the Spectral Partitioning algorithm into a complete function to factorise and solve large sparse linear systems using nested dissection is a primary research target, and would depend on a number of subsidiary goals, which are described below.

1. Problems in dealing with the choice of  $m$  in the Lanczos algorithm have been discussed in §4.3. Research is required to establish a more rigorous solution to this problem.
2. **decomp** should be augmented to become a function that also returns the **structure** of the permutation  $P$ . This might be made possible by using a ternary code associated with each element of  $P$ , to describe the level of recursion required to obtain it. This would facilitate a factorisation routine that recursively decomposed the large sparse matrix into manageable atomic-sized units. Implementation would be coupled with an analysis of the performance of this technique relative to other methods, such as Gauss-Seidel and Successive Over Relaxation.
3. Currently, there is currently very little error-checking in any of the functions and programs that I have written, but this facility is readily implemented in the context of the interface with MESCHACH. Input parameters are not carefully checked for existence and correctness of dimension, and there is not much checking through the codes for other problems, especially numerical problems. In part, the latter is due to time constraints, but it is also due to the total complexity of the algorithm. Further work is needed to make the routines more robust.
4. More efficient data-structures would improve programming and execution speed. In particular, the edge listing  $H$ , is currently represented as an  $|H| \times 2$  array of **doubles** (**MAT \*H**). (The same is true for the degree listing  $D$ .) As they are intended to be integers, this means that input, output and comparison operations on their elements require repeated casts. This is poor programming, and consumes more CPU time than required. A replacement system could involve simple integer arrays, or more powerfully, it could be a C structure containing an array, each element of which points to 2 integers representing adjacent edges, modelled after other MESCHACH structures. A possible data-structure and referencing description of this system is provided in Figure 1.

In conjunction with this we could develop a graph-input routine, together with a suite of error-checking routines. Currently, the graph is input as an adjacency matrix. This could be exchanged for an incidence matrix input. **decomp** could internally convert this to the appropriate adjacency matrix. This would improve the ability of the user to input a graph accurately. The driver program could be extended to handle a series of different input formats.

```

typedef struct edge_elt
{
    unsigned int    first, second;
} edge_elt

typedef struct edge_list
{
    unsigned int    m, max_m;
    ee              *edge_elt;
} edge_list;

PERM *decomp( ... )
{
    edge_list        *D, *H;
    /* "(H->ee[i]).first" now replaces "H->me[i][0]" */

```

Figure 1: C Data structure for edge-listing and degree-recording.

Other possibilities include storing the adjacency matrix of the graph as a dense **bit** array, rather than a sparse integer C structure. This would however, make implementation far more difficult, and is not considered to be practical for the internal workings of **decomp**. The programs **gr\_lap** and **testdc**, which involve input and output of graphs (as their adjacency matrices), could be made considerably more efficient by this practice.

5. The nested dissection algorithm lends itself to implementation on parallel processing machines, and a long-term goal could be to implement it in such an environment.
6. Pothen et al. [21] discuss various estimates for the size of an adequate separator set, calculable in terms of various eigenvalues of the graph Laplacian and other parameters of the graph. It might be valuable to investigate the use of these estimates for some internal consistency checks within the decomposition algorithm.
7. The implementation of a true minimum cover algorithm, as described in §2.5.1, would ensure that  $|S|$  is as small as possible. The algorithm presented in [18, pp221–227] appears to be the most efficient and accessible choice. It could be expected that the increase in computational expense in using this minimum cover algorithm would be compensated for by the reduction in computational expense (due to the smaller  $|S|$ ) in a full factorisation routine.

## 6 Listing of Graphic Examples

This appendix is a listing and discussion of some of the graphic examples used in developing the code and algorithms. Names and descriptions of the experimental graphs are presented in Table 2. Should these names appear unfamiliar, the reader is informed that they are loosely transliterated Hebrew names of various people in the Hebrew Bible. There is no particular significance in this choice of names.

Name	$n$	Description
Idit	7	Simplest case, almost a tree
Moshe	20	Hand-drawn example
Itzhak	31	Almost a tree
Shimuel	32	Simple example
Ishmail	55	$5 \times 11$ 5-point grid
Yacov	105	$5 \times 21$ 5-point grid
Yair	121	$11 \times 11$ 5-point grid
Arieh	505	$5 \times 101$ 5-point grid
Aaron	2121	$21 \times 101$ 5-point grid
Schlomo	6161	$61 \times 101$ 5-point grid
Shimshon	6400	$80 \times 80$ 5-point grid

Table 2: Listing of the experimental graphs.

The large grid graphs Schlomo ( $61 \times 101$ ) and Shimshon ( $80 \times 80$ ) are used for comparison with the results of [21].

### 6.1 Idit

Idit is the smallest graph examined (7 vertices, 8 edges), and is depicted in Figure 2. Its graph Laplacian is presented in Figure 3. Figures 4 and 5 are plots of the eigenvalue spectrum and the second eigenvector (respectively) of the graph Laplacian of Idit, as generated by MATLAB.<sup>4</sup>



Figure 2: Idit.

---

<sup>4</sup>MATLAB is an (interpreted) matrix computation package, and is a trademark of The Mathworks, Inc.

$$L = \begin{bmatrix} 2 & -1 & -1 & . & . & . & . \\ -1 & 2 & -1 & . & . & . & . \\ -1 & -1 & 3 & -1 & . & . & . \\ . & . & -1 & 2 & -1 & . & . \\ . & . & . & -1 & 3 & -1 & -1 \\ . & . & . & . & -1 & 2 & -1 \\ . & . & . & . & -1 & -1 & 2 \end{bmatrix}$$

Figure 3: The sparse graph Laplacian of Idit.

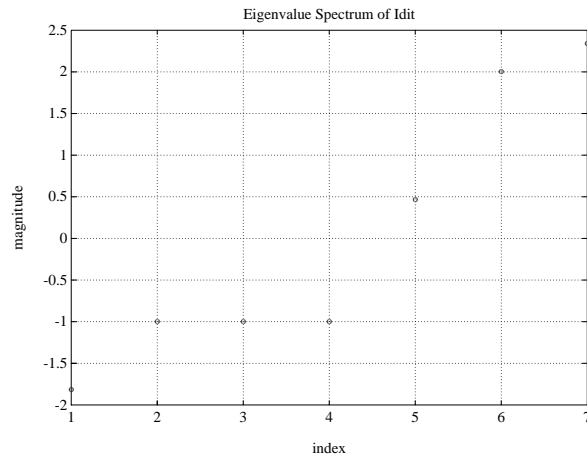


Figure 4: The eigenvalue spectrum of Idit.

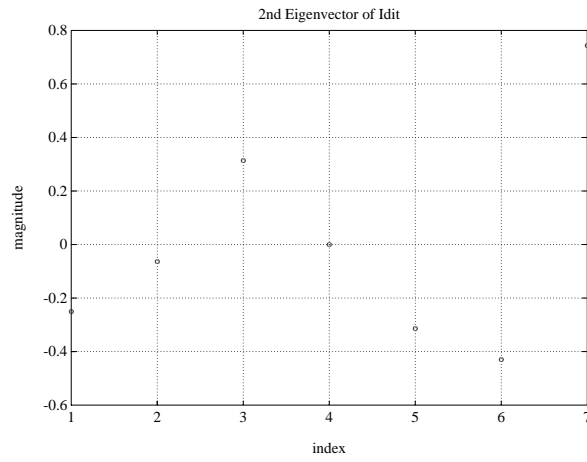


Figure 5: The second eigenvector of Idit.

## 6.2 Moshe

Moshe is a hand-drawn graph made for illustration of the progress of the Spectral Partitioning algorithm. It is a planar graph (20 vertices, 32 edges, diameter 7), and is depicted in Figure 6. The operation of the Spectral Partitioning algorithm (Algorithm 2.2) is described for Moshe in the following paragraphs.

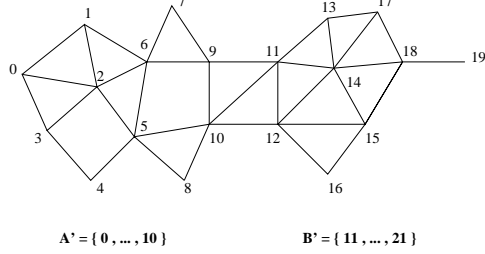


Figure 6: Moshe.

The initial graph partitioning process partitions  $N = \{0, \dots, 19\}$ , into 2 almost equal halves,  $A' = \{0, \dots, 9\}$  and  $B' = \{10, \dots, 19\}$ , with an **edge** separator set of size 3:  $E_1 = \{(8, 10), (9, 10), (9, 11)\}$ .  $A_1$  is the set of vertices in  $A'$  that are adjacent to vertices in  $B'$ , and vice-versa. Inspection shows that the relevant bipartite graph on  $E_1$  consists of  $A_1 = \{8, 9\}$  and  $B_1 = \{10, 11\}$ .

Using only  $A_1$ ,  $B_1$  and  $E_1$ , we seek to calculate a **vertex** separator set  $S$ , as a subset of the vertices in  $A_1$  and  $B_1$ , such that all edges in  $A_1$  and  $B_1$  are incident upon at least one vertex in  $S$ . We would like  $S$  to be of minimum cardinality. Whilst this may in general be a non-trivial problem, in this case inspection shows that  $\{9, 10\}$ ,  $\{8, 9\}$  and  $\{10, 11\}$  are minimum covers, the first case corresponding to Algorithm 2.3 accepting the initial choice of  $S$ , the others corresponding to the use of  $A_1$  or  $B_1$  as the cover.

Lastly, find  $A = A' \setminus S$  and  $B = B' \setminus S$ . Use of  $S = \{9, 10\}$  gives  $A = \{0, \dots, 8\}$  and  $B = \{11, \dots, 19\}$ . This choice is depicted in Figure 7. The resultant set of cut edges are drawn with dashed lines.

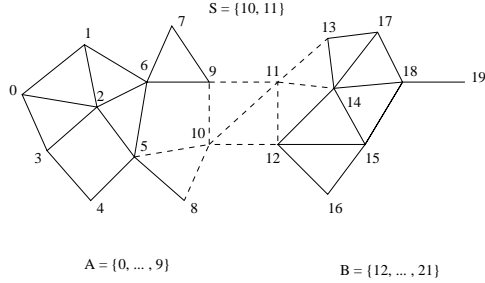


Figure 7: Moshe, after partitioning.

Figures 8 and 9 are plots of the eigenvalue spectrum and the second eigenvector (respectively) of the graph Laplacian of Moshe, as generated by MATLAB.

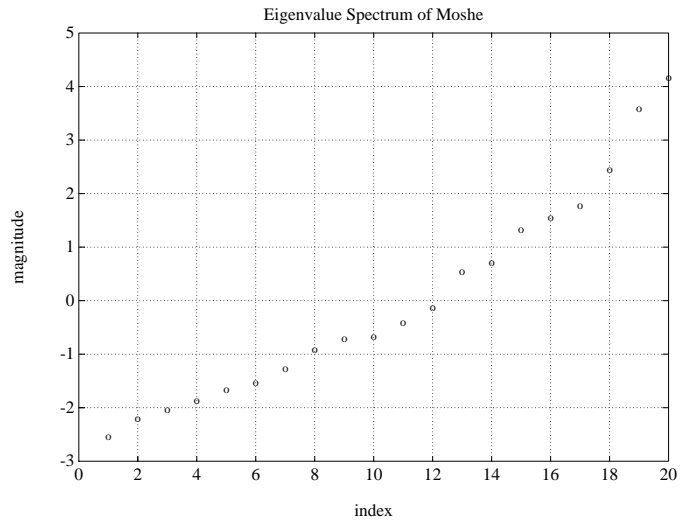


Figure 8: The eigenvalue spectrum of Moshe.



Figure 9: The second eigenvector of Moshe.

### 6.3 Itzchak

The graph Itzchak is presented in Figure 11. Its adjacency matrix is quite illustrative, and is presented in Figure 10. Figures 13 and 14 are plots of the eigenvalue spectrum and the second eigenvector (respectively) of the graph Laplacian of Itzchak, as generated by MATLAB.

1	1																
1		1															
1	1			1													
		1	1														
		1		1													
		1	1		1												
1			1					1									
				1	1												
			1		1												
			1	1			1										
					1	1		1									
					1			1									1
							1	1									
							1		1								
								1	1								
								1		1							1
									1	1							
										1	1						
											1	1				1	
												1	1			1	
													1	1			1
														1			
															1		
																1	
																	1

Figure 10: The adjacency matrix  $L$  of Itzchak.

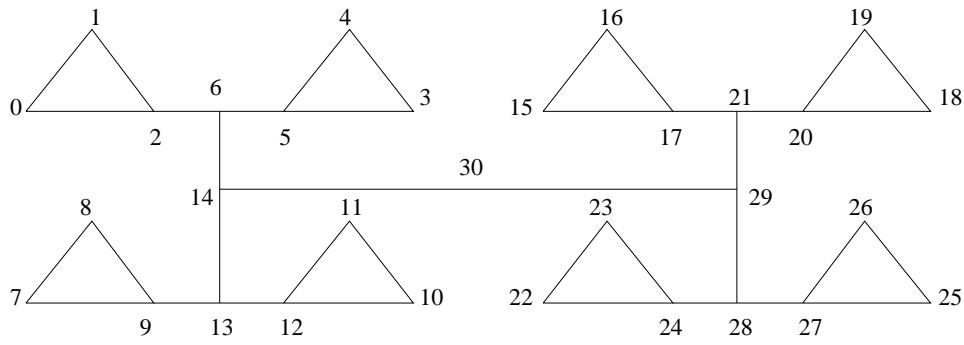


Figure 11: Itzhak.

## 6.4 Shimuel

Shimuel (Figure 12) is another example on a small number of vertices, also used for development purposes. It is set up to find the separator set  $S = \{11, 20\}$ . Figures 15 and 16 are plots of the eigenvalue spectrum and the second eigenvector (respectively) of the graph Laplacian of Shimuel, as generated by MATLAB.

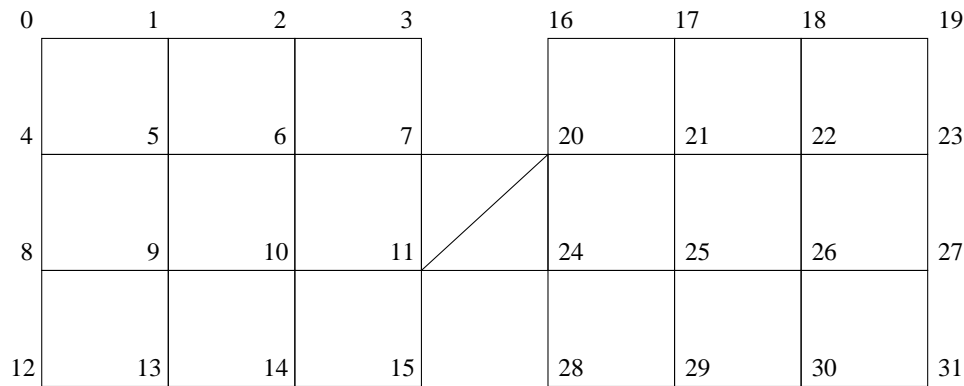


Figure 12: Shimuel.

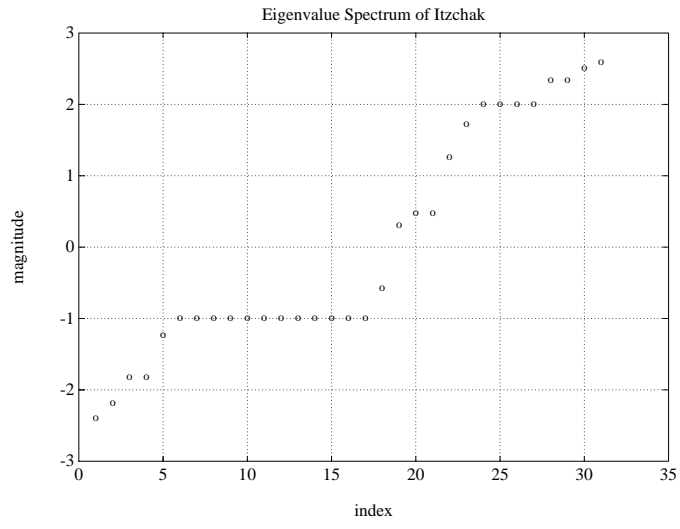


Figure 13: The eigenvalue spectrum of Itzchak.

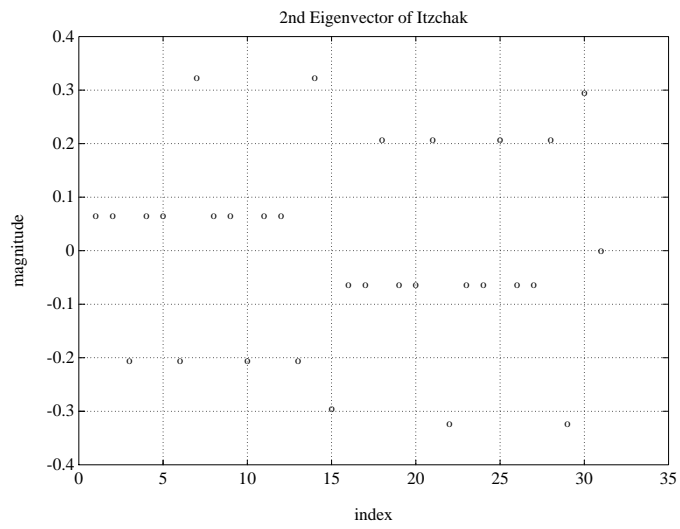


Figure 14: The second eigenvector of Itzchak.

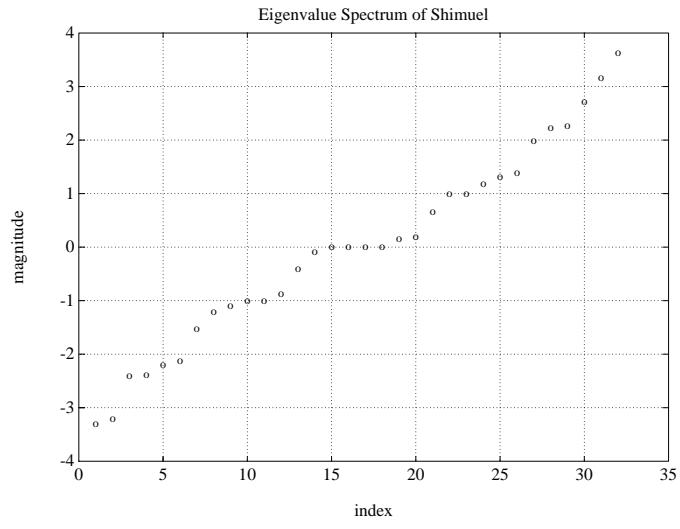


Figure 15: The eigenvalue spectrum of Shimuel.

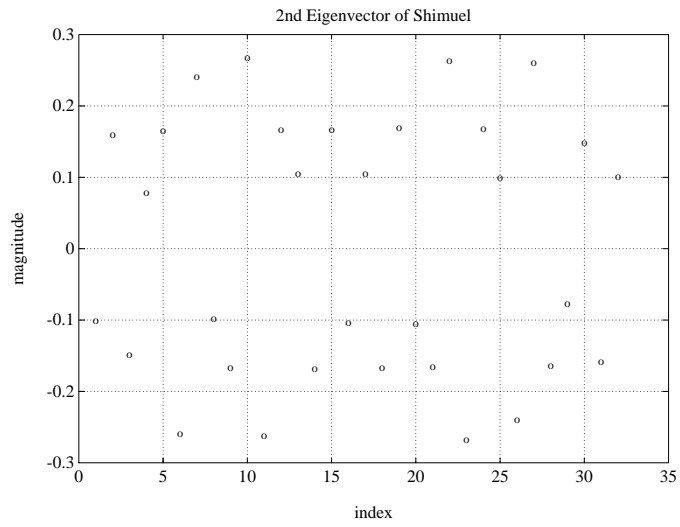


Figure 16: The second eigenvector of Shimuel.

## 6.5 Ishmail

Ishmail is the only (5-point) grid graph represented in this appendix, and is depicted in Figure 17. It is a  $5 \times 11$  grid ( $n = 55$  vertices). The adjacency matrix of Ishmail illustrates that referred to in §2.2, and its form is:

$$G = \begin{bmatrix} P & I_5 & . & . & . \\ I_5 & P & I_5 & . & . \\ . & I_5 & P & I_5 & . \\ . & . & I_5 & P & I_5 \\ . & . & . & I_5 & P \end{bmatrix},$$

where

$$P = \begin{bmatrix} . & 1 & . & . & . & . & . & . & . & . & . & . \\ 1 & . & 1 & . & . & . & . & . & . & . & . & . \\ . & 1 & . & 1 & . & . & . & . & . & . & . & . \\ . & . & 1 & . & 1 & . & . & . & . & . & . & . \\ . & . & . & 1 & . & 1 & . & . & . & . & . & . \\ . & . & . & . & 1 & . & 1 & . & . & . & . & . \\ . & . & . & . & . & 1 & . & 1 & . & . & . & . \\ . & . & . & . & . & . & 1 & . & 1 & . & . & . \\ . & . & . & . & . & . & . & 1 & . & 1 & . & . \\ . & . & . & . & . & . & . & . & 1 & . & 1 & . \\ . & . & . & . & . & . & . & . & . & 1 & . & . \end{bmatrix}.$$

	0	1	2	3	4	5	6	7	8	9	10	
11						16						21
22						27						32
33						38						43
44						49						54

Figure 17: Ishmail.

Figures 18 and 19 are plots of the eigenvalue spectrum and the second eigenvector (respectively) of the graph Laplacian of Ishmail, as generated by MATLAB.

Section 4 (titled “Graph Products”) of [21], discusses the expected repetition of the second eigenvector of an  $m \times n$  5-point grid graph, showing that it can be found as the Kronecker (tensor, outer) product of the length- $m$  path graph and an  $\mathbf{e}$ -vector of size  $n$ . Examination of Figure 19 (and Figure 21) clearly demonstrates this result. Clever exploitation of this property may reduce the overall computational expense required for the Spectral Partitioning algorithm, but not significantly, as the expense is dominated by the computations that yield  $\lambda_2$  (see §4.3). As this property is only true for **grid** graphs, implementation would require **decomp** to have an extra input flag.

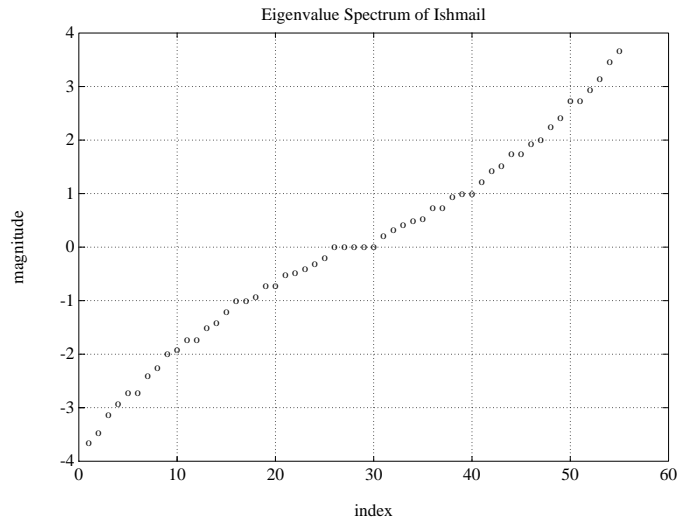


Figure 18: The eigenvalue spectrum of Ishmail.

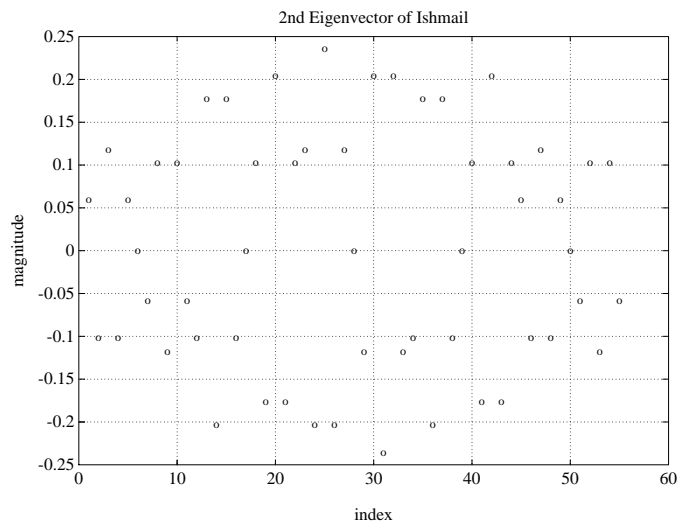


Figure 19: The second eigenvector of Ishmail.

## 6.6 Yacov

Yacov is a  $5 \times 21$  5-point grid graph on  $n = 105$  vertices, and is not depicted here. Figures 20 and 21 are plots of the eigenvalue spectrum and the second eigenvector (respectively) of the graph Laplacian of Yacov, as generated by MATLAB.

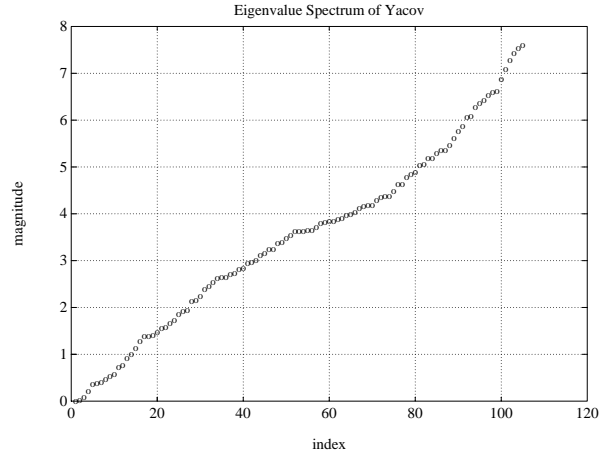


Figure 20: The eigenvalue spectrum of Yacov.

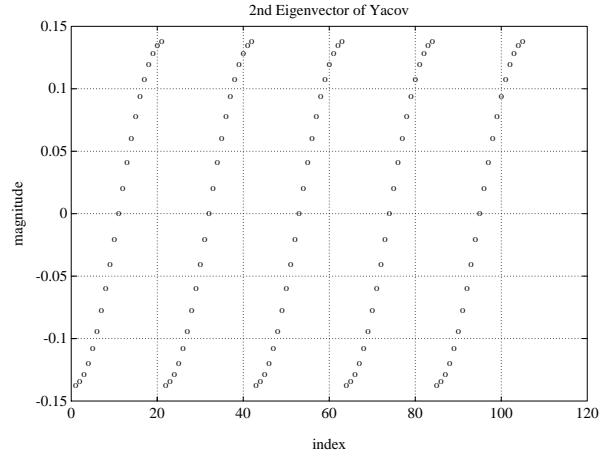


Figure 21: The second eigenvector of Yacov.

## 7 Code Listing

As mentioned in §3, the algorithms presented in this report are implemented in ANSI-standard C, and interface with MESCHACH, the C software library for numerical analysis written by David Stewart [23]. The contents of this appendix are listings of the source code that I have written for this work. Without access to [23], references to data structures and functions called from MESCHACH will be meaningless. Listings of the data-structures and functions are not provided here as they are quite long. §4.3 mentions that the code written is not necessarily bug-free, and currently has very little error-handling capacity.

### 7.1 Listing of decomp

```
/*                                decomp                                */

/* *****

Contents : This file is called "decomp.c", and contains the function
"decomp".

    Aim : (Recursively) decompose the vertex set of a (sparse) graph,
    represented by a graph Laplacian L, into separator sets.
    This function returns a pointer to a permutation, which,
    when applied to L, will decompose it for efficient
    recursive LDLT factorisation.

Language : ANSI Standard C

Author : David De Wit (and David Stewart) March 5 - June 19 1991

***** */

#include <math.h>
#include <stdio.h>
#include "matrix.h"
#include "matrix2.h"
#include "sparse.h"
#include "sparse2.h"

extern double select(VEC *, u_int);
extern int icmp(u_int *, u_int *);

u_int prt_tol, p;

/* *****

#define show_perm(PP) ((PP->size < prt_tol) ? out_perm(PP) :\
    printf("Permutation: size: %d\n", PP->size))
#define show_vec(VV) ((VV->dim < prt_tol) ? out_vec(VV) :\
    printf("Vector: dim: %d\n", VV->dim))
#define show_mat(MM) ((MM->m < prt_tol) ? out_mat(MM) :\
    printf("Matrix: m: %d by 2\n", MM->m))
```

```

PERM *decomp(sp_mat *L, PERM *P, PERM *A, PERM *B, PERM *S, int rec_lvl)
{
    int          i, j, k, l, n = L->m, tempA, tempB, tempH,
                inA, inB, inA1, inB1, inD, inS, inAdash, inBdash,
                N_nonzero_v, max_degree, vmax, imax, ev_good_enough;
    double        medval, L2old, L2new, L2tol = 1e-04, Rtol = 1e-01,
                yi, yj, sum, alpha, beta;
    PERM          *Adash, *Bdash, *A1, *B1, *AA, *AB, *AS, *BA, *BB, *BS,
                *pivot;
    VEC           *a, *b, *c, *d, *x0, *w, *y, *z, *resid, *V, *W, *TMP;
    MAT           *C, *I, *Q, *T, *D, *H;
    row_elt       *elt_list;
    sp_mat        *AL, *BL;

    if (!L || !A || !B || !S || !P)
        error(E_NULL, "decomp");
    if (L->m != L->n || P->size != n)
        error(E_SIZES, "decomp");

    printf("At top of decomp. rec_lvl: %d\n", rec_lvl);

/* 0. Set up all the required matrix and vector elements. Firstly, initialise
all those with dimensions fixed, depending on "n".*/

    Adash = get_perm(n);                Bdash = get_perm(n);
    A1 = get_perm(n);                  B1 = get_perm(n);
    y = get_vec(n);                    x0 = get_vec(n);
    resid = get_vec(n);                TMP = get_vec(n);
    V = get_vec(n);                   W = get_vec(n);
    D = get_mat(n*p, 2);               H = get_mat(n*p, 2);

/* Initialise size of variable-size data-structures 1. They are soon resized. */

    a = get_vec(1);                    b = get_vec(1);
    c = get_vec(1);                    d = get_vec(1);
    I = get_mat(1, 1);                 T = get_mat(1, 1);
    C = get_mat(1, 1);                 Q = get_mat(1, 1);
    w = get_vec(1);                    z = get_vec(1);
    pivot = get_perm(1);

/* 1. Use the Lanczos method and a tridiagonal eigendecomposition routine to
calculate the 2nd eigenvalue and corresponding eigenvector of L. Pothén, et al.
suggest the choice for the initial x0, and we begin with W = normalised(x0) and
V = LW. */
    for (i = 0; i < n; x0->ve[i] = i - (n - 1)/2, i++);
    L2old = 0;                          ev_good_enough = 0;
    while (!ev_good_enough)
    {
        j = 0;
        sv_mlt(1.0/n2(x0), x0, W);
        sp_mv_mlt(L, W, V);
        L2new = L2old + 2*L2tol;         beta = 1;
        while (j < 2 || fabs((L2old - L2new)/L2new) > L2tol)
        {
            j++;
            if (j*n > 250000)
            {
                printf("j.n = %d.%d = %d ", j, n, j*n);
                printf("Using too much memory, cutting out!\n");
                exit(0);
            }
            a = v_resize(a, j);          b = v_resize(b, j-1);
            if (j > 1)
                b->ve[j-2] = beta;
            Q = m_resize(Q, n, j);       set_col(Q, j-1, W);

/* Store W in Q. */
            a->ve[j-1] = alpha = in_prod(W, V);
            v_mltadd(V, W, -alpha, V);

```



```

/* 2. Calculate Adash, Bdash and H. Adash and Bdash are the vertex sets
generated by the partitioning of the vertices of G by H. H is an edge separator
set of G, found using Fiedler's method. y is the 2nd eigenvector of the
Laplacian matrix of G, with median value medval. */

/* 2.1. Calculate the median of the elements of the 2nd eigenvector. */
medval = select(y, (u_int) ((n + 1)/2));

/* 2.2.1. Set up Adash and Bdash. */

for (i = Adash->size = Bdash->size = 0; i < y->dim; i++)
    if (y->ve[i] <= medval)
        Adash->pe[Adash->size++] = i;
    else
        Bdash->pe[Bdash->size++] = i;

printf("First Pass Adash and Bdash made\n");
show_perm(Adash);
show_perm(Bdash);

/* 2.2. Set up H. Search through the upper half triangle of L, and, for
each edge encountered, insert it into H only if the eigenvaluation of
the vertices crosses the median. */

for (i = H->m = 0; i < L->m; i++)
{
    elt_list = (L->row[i]).elt;
    yi = y->ve[i];
    for (k = 0; k < (L->row[i]).len; k++)
        if ((j = elt_list[k].col) < i)
        {
            yj = y->ve[j];
            if (yi <= medval && yj > medval)
            {
                H->me[H->m][0] = i;
                H->me[H->m++][1] = j;
            }
            else if (yj <= medval && yi > medval)
            {
                H->me[H->m][0] = j;
                H->me[H->m++][1] = i;
            }
        }
}

printf("First Pass H made\n");
show_mat(H);

/* 2.2.2. If |Adash| - |Bdash| > 1, move enough vertices with components equal
to medval from Adash to Bdash. Not easy! Also must correct H. */

for (i = 0; i < Adash->size && Adash->size - Bdash->size > 1; i++)
    if (y->ve[i] == medval)
    {
        for (j = 0; Adash->pe[j] != i; j++);
        Adash->pe[j] = Adash->pe[--Adash->size];
        Bdash->pe[Bdash->size++] = i;
        for (k = 0; k < H->m; k++)
            for (l = 0; H->me[k][0] == i && l < 2; l++)
                H->me[k][l] = H->me[H->m--][l];
    }
freevec(y);

/* Sort Adash and Bdash. */
qsort(Adash->pe, Adash->size, sizeof(u_int), icmp);
qsort(Bdash->pe, Bdash->size, sizeof(u_int), icmp);

printf("2nd Pass Adash: ");
printf("\n2nd Pass Bdash: ");
printf("\n2nd Pass H: ");

show_perm(Adash);
show_perm(Bdash);
show_mat(H);

```

/\* 3. Calculate S, a vertex separator set of G. A reminder that H is an edge separator set, and Adash and Bdash are a disjoint cover of the vertex set of G. A1 and B1 are the respective subsets of Adash and Bdash that consist of only vertices that adjoin edges in H. The two vertex sets A and B are found, such that  $A = \text{Adash} \setminus S$ ,  $B = \text{Bdash} \setminus S$ . Respective edges are not relevant to the further progress of "decomp", and so are not found. S is optimally as small as possible, and one is found that is hopefully relatively small, by the better of 2 algorithms that seem like a quick and good way:

1. Collect vertices of reducing degree from A1 U B1 until H is covered.
2. Use all the vertices of the smaller of the two sets, A1 and B1. \*/

/\* 3.1.1. Set up D, a listing of the degrees of vertices in H. \*/

```
for (i = D->m = 0; i < H->m; i++)
    for (j = 0; j < 2; j++)
    {
        tempH = (u_int) H->me[i][j];
        for (k = inD = 0; k < D->m; k++)
            if ((u_int) D->me[k][0] == tempH)
                D->me[k][inD = 1]++;

        if (!inD)
        {
            D->me[D->m][0] = tempH;
            D->me[D->m][1]++;
        }
    }

printf("\nD: ");
show_mat(D);
```

/\* 3.1.2. Set up A1 and B1, respective sides of H. \*/

```
for (i = A1->size = B1->size = 0; i < H->m; i++)
{
    tempH = H->me[i][0];
    for (k = inA1 = 0; k < A1->size; k++)
        inA1 = inA1 || A1->pe[k] == tempH;
    if (!inA1)
        A1->pe[A1->size++] = tempH;
    tempH = H->me[i][1];
    for (k = inB1 = 0; k < B1->size; k++)
        inB1 = inB1 || B1->pe[k] == tempH;
    if (!inB1)
        B1->pe[B1->size++] = tempH;
}
```

/\* Sort A1 and B1. Not currently useful, but good for the future. \*/

```
qsort(A1->pe, A1->size, sizeof(u_int), icmp);
qsort(B1->pe, B1->size, sizeof(u_int), icmp);

printf("2nd Pass A1: ");
show_perm(A1);
printf("\n2nd Pass B1: ");
show_perm(B1);
```

/\* 3.2. Establish the permutation of the vertex separator set S, via a grungy but apparently working piece of code. The algorithm used:

1. Find the vertex of highest degree in D, such that:  
D->me[imax][0] = vmax; D->me[imax][1] = max\_degree;
2. Set the degree of this vertex to zero, decrement the number of non-zero vertices, and add this vertex to S.  
D->me[imax][1] = 0; N\_nonzero\_v--; S->pe[S->size++] = vmax;
3. Search out the edges in H containing this vertex, and decrement the vertices adjacent to it in D.
4. Repeat until no non-zero degree vertices remain. \*/

```
N_nonzero_v = D->m;
while (N_nonzero_v > 0)
{
    S->size = 0;

    for (i = max_degree = 0; i < D->m; i++)
        if (D->me[i][1] > max_degree)
            max_degree = (u_int) D->me[imax = i][1];

    D->me[imax][1] = 0;
    N_nonzero_v--;
    S->pe[S->size++] = vmax = (u_int) D->me[imax][0];
}
```

```

        for (i = 0; i < H->m && N_nonzero_v > 0 && max_degree > 0; i++)
            for (j = 0; j < 2 && N_nonzero_v > 0 && max_degree > 0; j++)
                if ((u_int) H->me[i][j] == vmax)
                {
                    k = (u_int) H->me[i][!j];
                    for (l = 0; l < D->m && (u_int) D->me[l][0] != k; l++);
                    if (D->me[l][1] > 0)
                    {
                        D->me[l][1]--;
                        max_degree--;
                        if ((u_int) D->me[l][1] == 0)
                            N_nonzero_v--;
                    }
                }
    }
}

/* Sort S. */
qsort(S->pe, S->size, sizeof(u_int), icmp);

printf("\nFirst pass S: ");          show_perm(S);

/* 3.3. Test to see if the S that has been created is larger than the
smaller of A1 and B1. If it is, replace S with this smaller set. */

if (S->size > A1->size)
{
    S = cp_perm(A1, S);              S->size = A1->size;
}
if (S->size > B1->size)
{
    S = cp_perm(B1, S);              S->size = B1->size;
}
printf("\nSecond pass S: ");        show_perm(S);

/* 3.4. Create A, B, such that A, B and S are a disjoint cover of N. */
for (i = A->size = 0; i < Adash->size; i++)
{
    k = Adash->pe[i];
    for (j = inS = 0; j < S->size; j++)
        inS = inS || S->pe[j] == k;
    if (!inS)
        A->pe[A->size++] = k;
}
for (i = B->size = 0; i < Bdash->size; i++)
{
    k = Bdash->pe[i];
    for (j = inS = 0; j < S->size; j++)
        inS = inS || S->pe[j] == k;
    if (!inS)
        B->pe[B->size++] = k;
}

/* Sort A and B. */
qsort(A->pe, A->size, sizeof(u_int), icmp);
qsort(B->pe, B->size, sizeof(u_int), icmp);

/* Print out the vital statistics of the whole process. */
printf("\nSizes of various elements:\n\n");
printf("\t# Adash, Bdash:   %d\t%d\n", Adash->size, Bdash->size);
printf("\tLambda2:          %f\n", L2new);
printf("\t# D:              %d\n", D->m);
printf("\t# H:              %d\n", H->m);
printf("\t# A1 and B1:       %d\t%d\n", A1->size, B1->size);
printf("A: ");              show_perm(A);
printf("B: ");              show_perm(B);
printf("S: ");              show_perm(S);
printf("*****\n\n");

/* Dump unneeded data-structures here! */
freeperm(A1);                freeperm(Adash);    freemat(D);
freeperm(B1);                freeperm(Bdash);    freemat(H);

```

```

/* 4. Recursively partition the sets A and B by calling "decomp" on them. */

/* 4.1. Decompose A and B if their sizes are > 3. For example, for A, a sparse
matrix AL is created (found by extracting the parts of L with indices in A),
then decomp is called with AL as L, A as P, and dummies AA, AB, AS as the other
parameters. In fact, the routine does not actually need these other parameters,
but they are used, as they may be required for further programming to also
return the structure of the final permutation. Required structures are created
in situ. */
    if (A->size > 3 && rec_lvl != -1)
    {
        AA = get_perm(A->size);
        AB = get_perm(A->size);
        AS = get_perm(A->size);
        AL = sp_get_mat(A->size, A->size, 3*p);

/* The following 9 lines of _slow_ code (to set the entries of AL) are a waste
of time, but I haven't yet figured out what else to do. Ditto for BL. In fact,
it appears that this is the slowest part of the whole algorithm! */

        for (i = 0; i < A->size; i++)
            for (j = 0; j < A->size; j++)
            {
                k = (u_int) sp_get_val(L, A->pe[i], A->pe[j]);
                if (k != 0)
                    sp_set_val(AL, i, j, k);
            }
        for (i = 0; i < A->size; i++)
            sp_set_val(AL, i, i, (AL->row[i]).len - 1);

/* *****
The alternative version uses the elements of A as sorted into increasing order.
Now we can scan through the rows of L that corresponded to entries in A, and
know that we only have to check the entries in the row of L up until the column
number becomes >= row number. It looks something like:

        for (i = 0; i < A->size; i++)
        {
            k = A->pe[i];
            elt_list = (L->row[k]).elt;
            for (j = 0; j < (L->row[k]).len; j++)
            {
                for (l = inA = 0; l < A->size; l++)
                    inA = inA || A->pe[l] == (elt_list[j]).val;
                if (inA)
                    sp_set_val(AL, i, j, -1);
            }
        }
        for (i = 0; i < A->size; i++)
            sp_set_val(AL, i, i, (AL->row[i]).len);
***** */

        if (!ck_symm(AL))
        {
            printf("Quitting as AL not symmetric!\n");
            exit(0);
        }
        if (!ck_sums(AL))
        {
            printf("Quitting as AL has invalid row sums!\n");
            exit(0);
        }
        decomp(AL, A, AA, AB, AS, ++rec_lvl);      rec_lvl--;
        freeperm(AA); freeperm(AB); freeperm(AS); sp_free_mat(AL);
    }
}

```

```

    if (B->size > 3 && rec_lvl != -1)
    {
        BA = get_perm(B->size);
        BB = get_perm(B->size);
        BS = get_perm(B->size);
        BL = sp_get_mat(B->size, B->size, 3*p);

        for (i = 0; i < B->size; i++)
            for (j = 0; j < B->size; j++)
            {
                k = (u_int) sp_get_val(L, B->pe[i], B->pe[j]);
                if (k != 0)
                    sp_set_val(BL, i, j, k);
            }
        for (i = 0; i < B->size; i++)
            sp_set_val(BL, i, i, (BL->row[i]).len - 1);

        if (!ck_symb(BL))
        {
            printf("Quitting as BL not symmetric!\n");
            exit(0);
        }
        if (!ck_sums(BL))
        {
            printf("Quitting as BL has invalid row sums!\n");
            exit(0);
        }
        decomp(BL, B, BA, BB, BS, ++rec_lvl);          rec_lvl--;
        freeperm(BA); freeperm(BB); freeperm(BS);    sp_free_mat(BL);
    }

/* 4.2. Map the apparent A, B, S to the actual A, B, S, using P, the permutation
of their actual names. */
    for (i = 0; i < A->size; i++)
        A->pe[i] = P->pe[A->pe[i]];
    for (i = 0; i < B->size; i++)
        B->pe[i] = P->pe[B->pe[i]];
    for (i = 0; i < S->size; i++)
        S->pe[i] = P->pe[S->pe[i]];

    for (i = 0; i < A->size; i++)
        P->pe[i] = A->pe[i];
    j = i;
    for (i = 0; i < B->size; i++)
        P->pe[i + j] = B->pe[i];
    j += i;
    for (i = 0; i < S->size; i++)
        P->pe[i + j] = S->pe[i];

    return P;
}

/* ***** */

int icmp(u_int *p1, u_int *p2)
{
    return *p1 - *p2;
}

```

## 7.2 Listing of mk\_sp\_graph

```
/*                                mk_sp_graph                                */

/* *****

Contents : This file is called "mk_sp_graph.c", and contains the
          function "mk_sp_graph".

          Aim : Generate a random G, the sparse matrix of a (sparse)
          undirected, unvaluated graph on n vertices with an average
          degree of p (number of edges/vertex). Further, the graph is
          almost planar, and this is achieved by using only a narrow
          bandwidth (q) in the original matrix, followed by a
          scrambling of its indices.

Language : ANSI Standard C

Author  : David De Wit   March 5 - May 27 1991

***** */

#include <stdio.h>
#include "matrix.h"
#include "sparse.h"

sp_mat *mk_sp_graph(sp_mat *G, u_int p, u_int q)
{
    u_int      i, j, n = G->n;
    double     temp, limit;

    srand(4);
    limit = ((double) p)*2147483648/((double) n);
    for (i = 0; i < n; i++)
        for (j = i + 1; (j < n) && (j < i + q); j++)
            if (rand() < limit)
            {
                temp = sp_set_val(G, i, j, 1.0);
                temp = sp_set_val(G, j, i, 1.0);
            }

    return G;
}
```

## 7.3 Listing of select

```
/*                                select                                */

/* *****

Contents : This file is called "select.c", and contains the function
"select".

    Aim : Select the kth smallest entry in vector a. On exit, the
    desired element is in its correct place. In particular,
    calling select(a, int((a->dim + 1)/2)) finds the median.

    See "Algorithms", Sedgewick (1983), p128. QA76.6.S435

Language : ANSI Standard C

Author : David De Wit    March 4 - May 27 1991

***** */

#include <stdio.h>
#include "matrix.h"

double select(VEC *a, u_int k)
{
    int          i, j, l, r;
    double        t, v;
    VEC           *b;

    b = get_vec(a->dim);          b = cp_vec(a, b);
    l = 0;                       r = b->dim - 1;
    while (r > l)
    {
        v = b->ve[r];             i = l - 1;      j = r;
        do
        {
            for (i++; b->ve[i] < v; i++);
            for (j--; b->ve[j] > v; j--);
            t = b->ve[i];
            b->ve[i] = b->ve[j];
            b->ve[j] = t;
        } while (j > i);
        b->ve[j] = b->ve[i];
        b->ve[i] = b->ve[r];
        b->ve[r] = t;
        if (i >= k)
            r = i - 1;
        if (i <= k)
            l = i + 1;
    }
    return (b->ve[k - 1]);
}
```

## 7.4 Listing of gr\_lap

```
/*                                gr_lap                                */

/* *****

Contents : This file is called "gr_lap.c", and contains the program of
the same name.

    Aim : Make the graph for solving Laplace's Equation on an M x N
    grid.

Language : ANSI Standard C

Author  : David De Wit (and David Stewart) May 8 - May 27 1991

***** */

#include <stdio.h>
#include <strings.h>
#include "matrix.h"
#include "sparse.h"

#define index(i,j) (N*((i)-1)+(j)-1)

main(int argc, char *argv[])
{
    u_int      i, j, M = atoi(argv[1]), N = atoi(argv[2]);
    char       *outname;
    sp_mat     *A;

    A = sp_get_mat(M*N, M*N, 5);
    for (i = 1; i <= M; i++)
        for (j = 1; j <= N; j++)
        {
            if (i < M)
                sp_set_val(A, index(i,j), index(i+1,j), 1);
            if (i > 1)
                sp_set_val(A, index(i,j), index(i-1,j), 1);
            if (j < N)
                sp_set_val(A, index(i,j), index(i,j+1), 1);
            if (j > 1)
                sp_set_val(A, index(i,j), index(i,j-1), 1);
        }
    outname = strcat(strcat(strcat("Lap.", argv[1]), "."), argv[2]);
    sp_fout_mat(fopen(outname, "w"), A);
}
```

## 7.5 Listing of testdc

```
/*                                testdc                                */

/* *****

Contents : This file is called "testdc.c", and contains the program of
the same name. It calls the function "decomp".

        Aim : Test the function "decomp" by setting-up and solving a
        problem.

Language : ANSI Standard C

Author  : David De Wit    March 5 - June 19 1991

***** */

#include <stdio.h>
#include "matrix.h"
#include "matrix2.h"
#include "sparse.h"

extern u_int    prt_tol, p, q;
extern PERM     *decomp(sp_mat *, PERM *, PERM *, PERM *, PERM *, int);
extern sp_mat   *mk_sp_graph(sp_mat *, u_int, u_int);
extern int      ck_symm(sp_mat *), ck_sums(sp_mat *);

main(int argc, char *argv[])
{
    u_int        i, j, k, n = atoi(argv[1]), q = 6, nfiles = 11;
    int          j_idx, rec_lvl = atoi(argv[4]);
    PERM         *A, *B, *S, *P;
    FILE         *fp;
    row_elt      *e;
    sp_row       *r;
    sp_mat       *G;
    char         *fname[] = {"Idit", "Moshe", "Itzchak", "Shimuel",
                             "Ishmail", "Yacov", "Yair", "Arieh",
                             "Aaron", "Schlomo", "Shimshon"};

    /* 0. Initialise some structures and constants for the problem. */
    prt_tol = atoi(argv[3]);          p = atoi(argv[2]);
    A = get_perm(n);                  B = get_perm(n);
    S = get_perm(n);                  P = get_perm(n);
    P = px_id(P);
```

```

/* 1. Set up problem by randomly generating or reading in a sparse matrix. */
printf("\nEnter a choice for the initial graph:\n\n");
printf("\t0: Make a random sparse graph\n");
for (i = 0; i < nfiles; i++)
    printf("\t%d: Read file \"%s\"\n", i + 1, fname[i]);
printf("\nYour Choice:\n");
scanf("%d", &i);
if (i)
{
    printf("\nReading a sparse graph from \"%s\"\n", fname[i - 1]);
    fp = fopen(fname[i - 1], "r");
    G = sp_fin_mat(fp);
}
else
{
    printf("\nMaking a sparse graph on %d vertices\n", n);
    printf("at an average %d edges/vertex.\n", p);
    G = sp_get_mat(n, n, 3*p);
    G = mk_sp_graph(G, p, q);
}

if (G->m != n)
    error(E_SIZES, "testdc");

/* 2. Calculate the sparse matrix of the graph Laplacian (L) of the graph
represented by the sparse matrix G. Defining d[i] as the degree of vertex i in
G; then L = diag(d) - G. All row and column sums in L = 0. G overwrites L. */

/* Ensure diagonal entries are in rows. */
for (i = 0; i < G->m; sp_set_val(G, i, i, 1.0), i++);

/* Set the values of the entries. */
for (i = 0; i < G->m; i++)
{
    r = &(G->row[i]);
    /* scan entries of row r */
    for (j_idx = 0, e = r->elt; j_idx < r->len; j_idx++, e++)
        if (e->col == i)
            e->val = r->len - 1; /* diagonal entry */
        else
            e->val = -1.0; /* off-diagonal entry */
}

if (!ck_symm(G))
{
    printf("Quitting as graph Laplacian not symmetric!\n");
    exit(0);
}
if (!ck_sums(G))
{
    printf("Quitting as graph Laplacian has invalid row sums!\n");
    exit(0);
}

/* 3. Call the function to do the partitioning. */
printf("\nCalling decomp from testdc ...\n\n");
P = decomp(G, P, A, B, S, rec_lvl);
printf("\nBack in Kansas ... P:\n");
if (P->size < prt_tol)
    out_perm(P);
else
    printf("Permutation, size: %d\n", P->size);
}

```

```

/* ***** */

int ck_symm(sp_mat *L)
{
    int i;
    static VEC *x=NULL, *y1=NULL, *y2=NULL;

    if (!L)
        error(E_NULL,"ck_symm");
    if (L->m != L->n)
        return FALSE;

    x = v_resize(x, L->m);          y1 = v_resize(y1, L->m);
    y2 = v_resize(y2, L->m);

    for (i = 0; i < 3; i++)
    {
        rand_vec(x);                y1 = sp_mv_mlt(L, x, y1);
        y2 = sp_vm_mlt(L, x, y2);    y1 = v_sub(y1, y2, y1);
        if (n2(y1) > L->m*MACHEPS)
            return FALSE;
    }

    return TRUE;
}

/* ***** */

int ck_sums(sp_mat *L)
{
    int i, j;
    double sum;
    sp_row *r;

    if (!L)
        error(E_NULL,"ck_sums");
    for (i = 0; i < L->m; i++)
    {
        r = &(L->row[i]);
        for (j = sum = 0; j < r->len; j++)
            sum += r->elt[j].val;
        if (fabs(sum) > L->m*MACHEPS)
            return FALSE;
    }

    return TRUE;
}

```

## References

- [1] W N Anderson Jr and T D Morley. Eigenvalues of the Laplacian of a graph. Technical Report TR-41-45, University of Maryland, October 1971.
- [2] Abraham Berman and Robert J Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, New York, 1979.
- [3] Iain S Duff, Roger G Grimes, and John Gregg Lewis. Sparse matrix test problems. *ACM Transactions on Mathematical Software*, 15(1):1–14, March 1989.
- [4] A L Dulmage and N S Mendelsohn. Two algorithms for bipartite graphs. *Journal of the Society for Industrial and Applied Mathematics*, 11:183–194, 1963.
- [5] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.
- [6] Miroslav Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory. *Czechoslovak Mathematical Journal*, 25(100):619–633, 1975.
- [7] Miroslav Fiedler and Vlastimil Pták. On matrices with non-positive off-diagonal elements and positive principal minors. *Czechoslovak Mathematical Journal*, 22(87):382–400, 1962.
- [8] J A George. Nested dissection of a regular finite element mesh. *SIAM Journal on Numerical Analysis*, 10:345–363, 1973.
- [9] Gene H Golub and Charles F Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 2nd edition, 1989.
- [10] D M Johnson, A L Dulmage, and N S Mendelsohn. Connectivity and reducibility of graphs. *Canadian Journal of Mathematics*, 14:529–539, 1962.
- [11] B W Kernighan and S Lin. An efficient heuristic procedure for partitioning graphs. *Bell Systems Technical Journal*, 49:291–307, 1970.
- [12] C Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research (of the National Bureau of Standards)*, 45:255–282, 1950.
- [13] C E Leiserson and John Gregg Lewis. Orderings for parallel sparse symmetric factorisation. In *Third SIAM Conference on Parallel Processing for Scientific Computing*. SIAM Publications, 1987. Referenced in [21].
- [14] Joseph W H Liu. A graph partitioning algorithm by node separators. *ACM Transactions on Mathematical Software*, 15(3):198–219, September 1989.
- [15] Henryk Minc. *Nonnegative Matrices*. Wiley, New York, 1988.
- [16] Robert Z Norman and Michael O Rabin. An algorithm for a minimum cover of a graph. *Proceedings of the American Mathematical Society*, 10:315–319, 1959.

- [17] A Ostrowski. Über die determinanten mit überwiegender hauptdiagonale und die absolute konvergenz von linearen iterationsprossen. *Commentarii Mathematici Helvetici*, 10:69–96, 1937. Referenced in [19].
- [18] Christos H Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice Hall, Englewood Cliffs, NJ, USA, 1982.
- [19] George Poole and Thomas Boullion. A survey on M-matrices. *SIAM Review*, 16:419–427, 1974.
- [20] Alex Pothén and Chin-Ju Fan. Computing the block triangular form of a sparse matrix. *ACM Transactions on Mathematical Software*, 16(4):303–324, 1990.
- [21] Alex Pothén, Horst D Simon, and Kang-Pu Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM Journal on Matrix Analysis and Applications*, 11(3):430–452, 1990.
- [22] Robert Sedgewick. *Algorithms*. Addison-Wesley, Reading, Massachusetts, 2nd edition, 1988.
- [23] David E Stewart. MESCHACH: Matrix Computations in C. Unpublished manuscript, 1991.
- [24] Richard S Varga. *Matrix Iterative Analysis*. Prentice Hall, Englewood Cliffs, 1962.